Brookhaven Scientific Computing

A rapidly changing environment

Eric Lançon

August 31, 2016

BROOKHAVEN NATIONAL LABORATORY

a passion for discovery

Office of Science
U.S. DEPARTMENT OF ENERGY

The RHIC Accelerator Complex

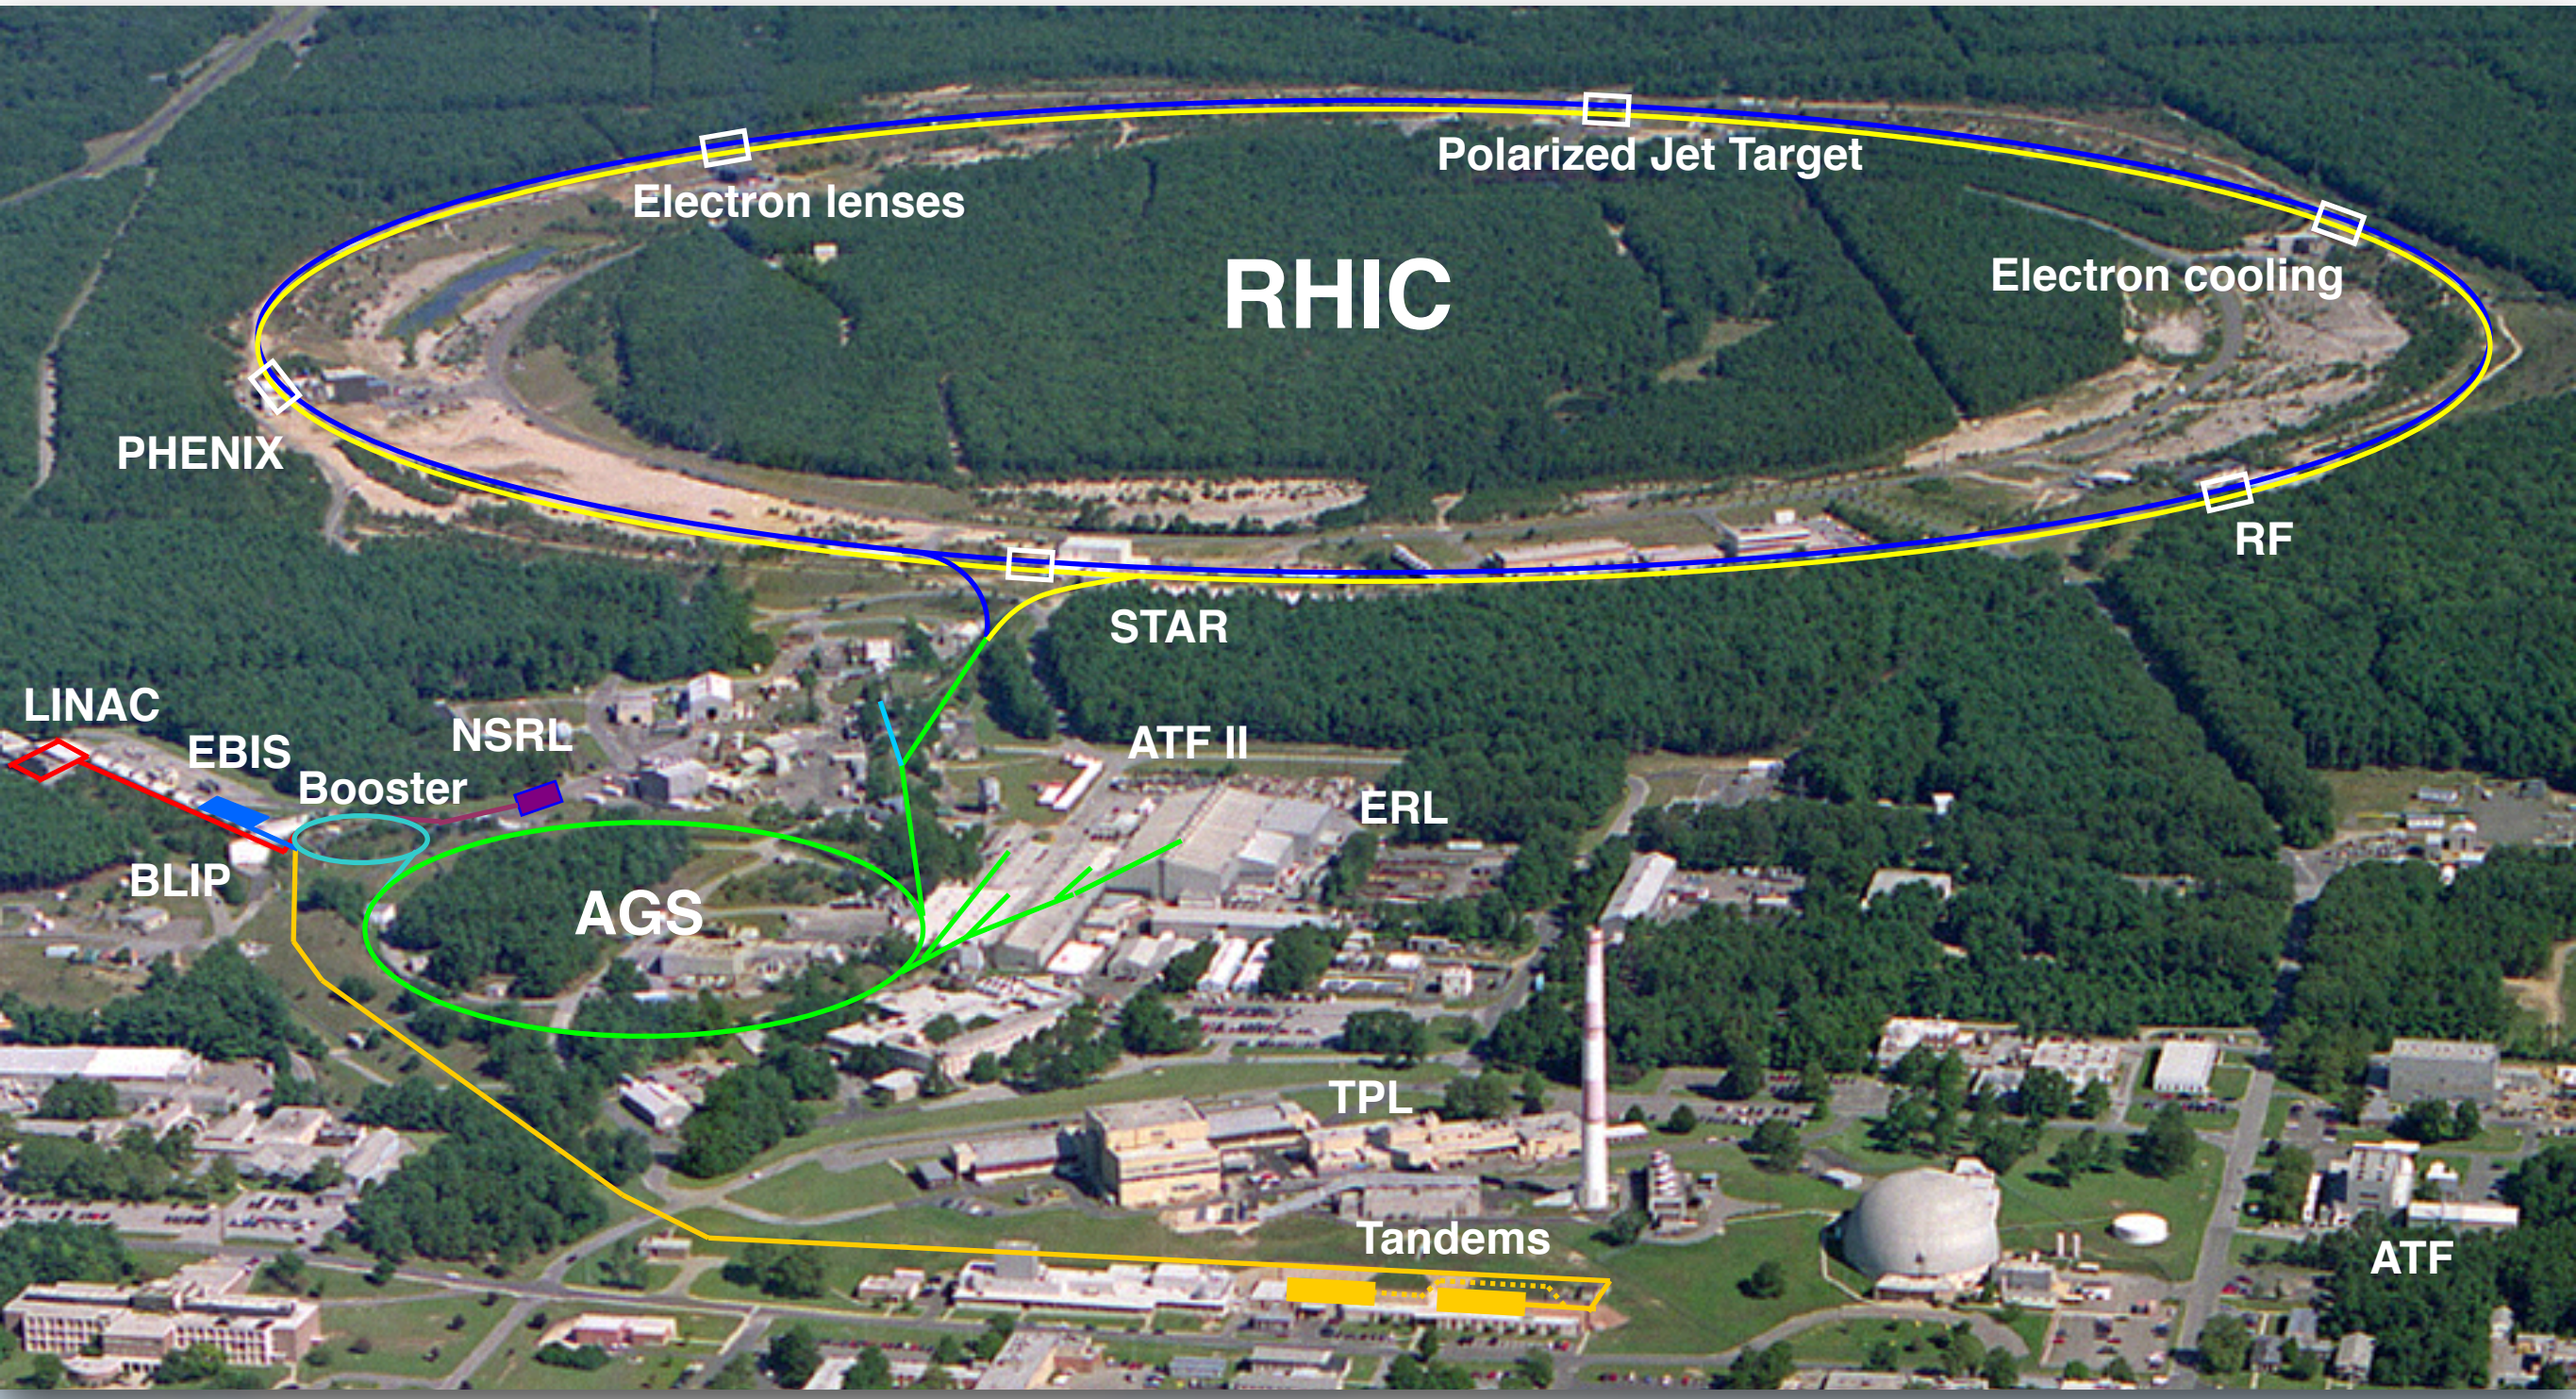

- Highly flexible and only US Hadron Collider
- Injectors also provide beams for unique applications

RHIC: 2 main experiments PHENIX & STAR

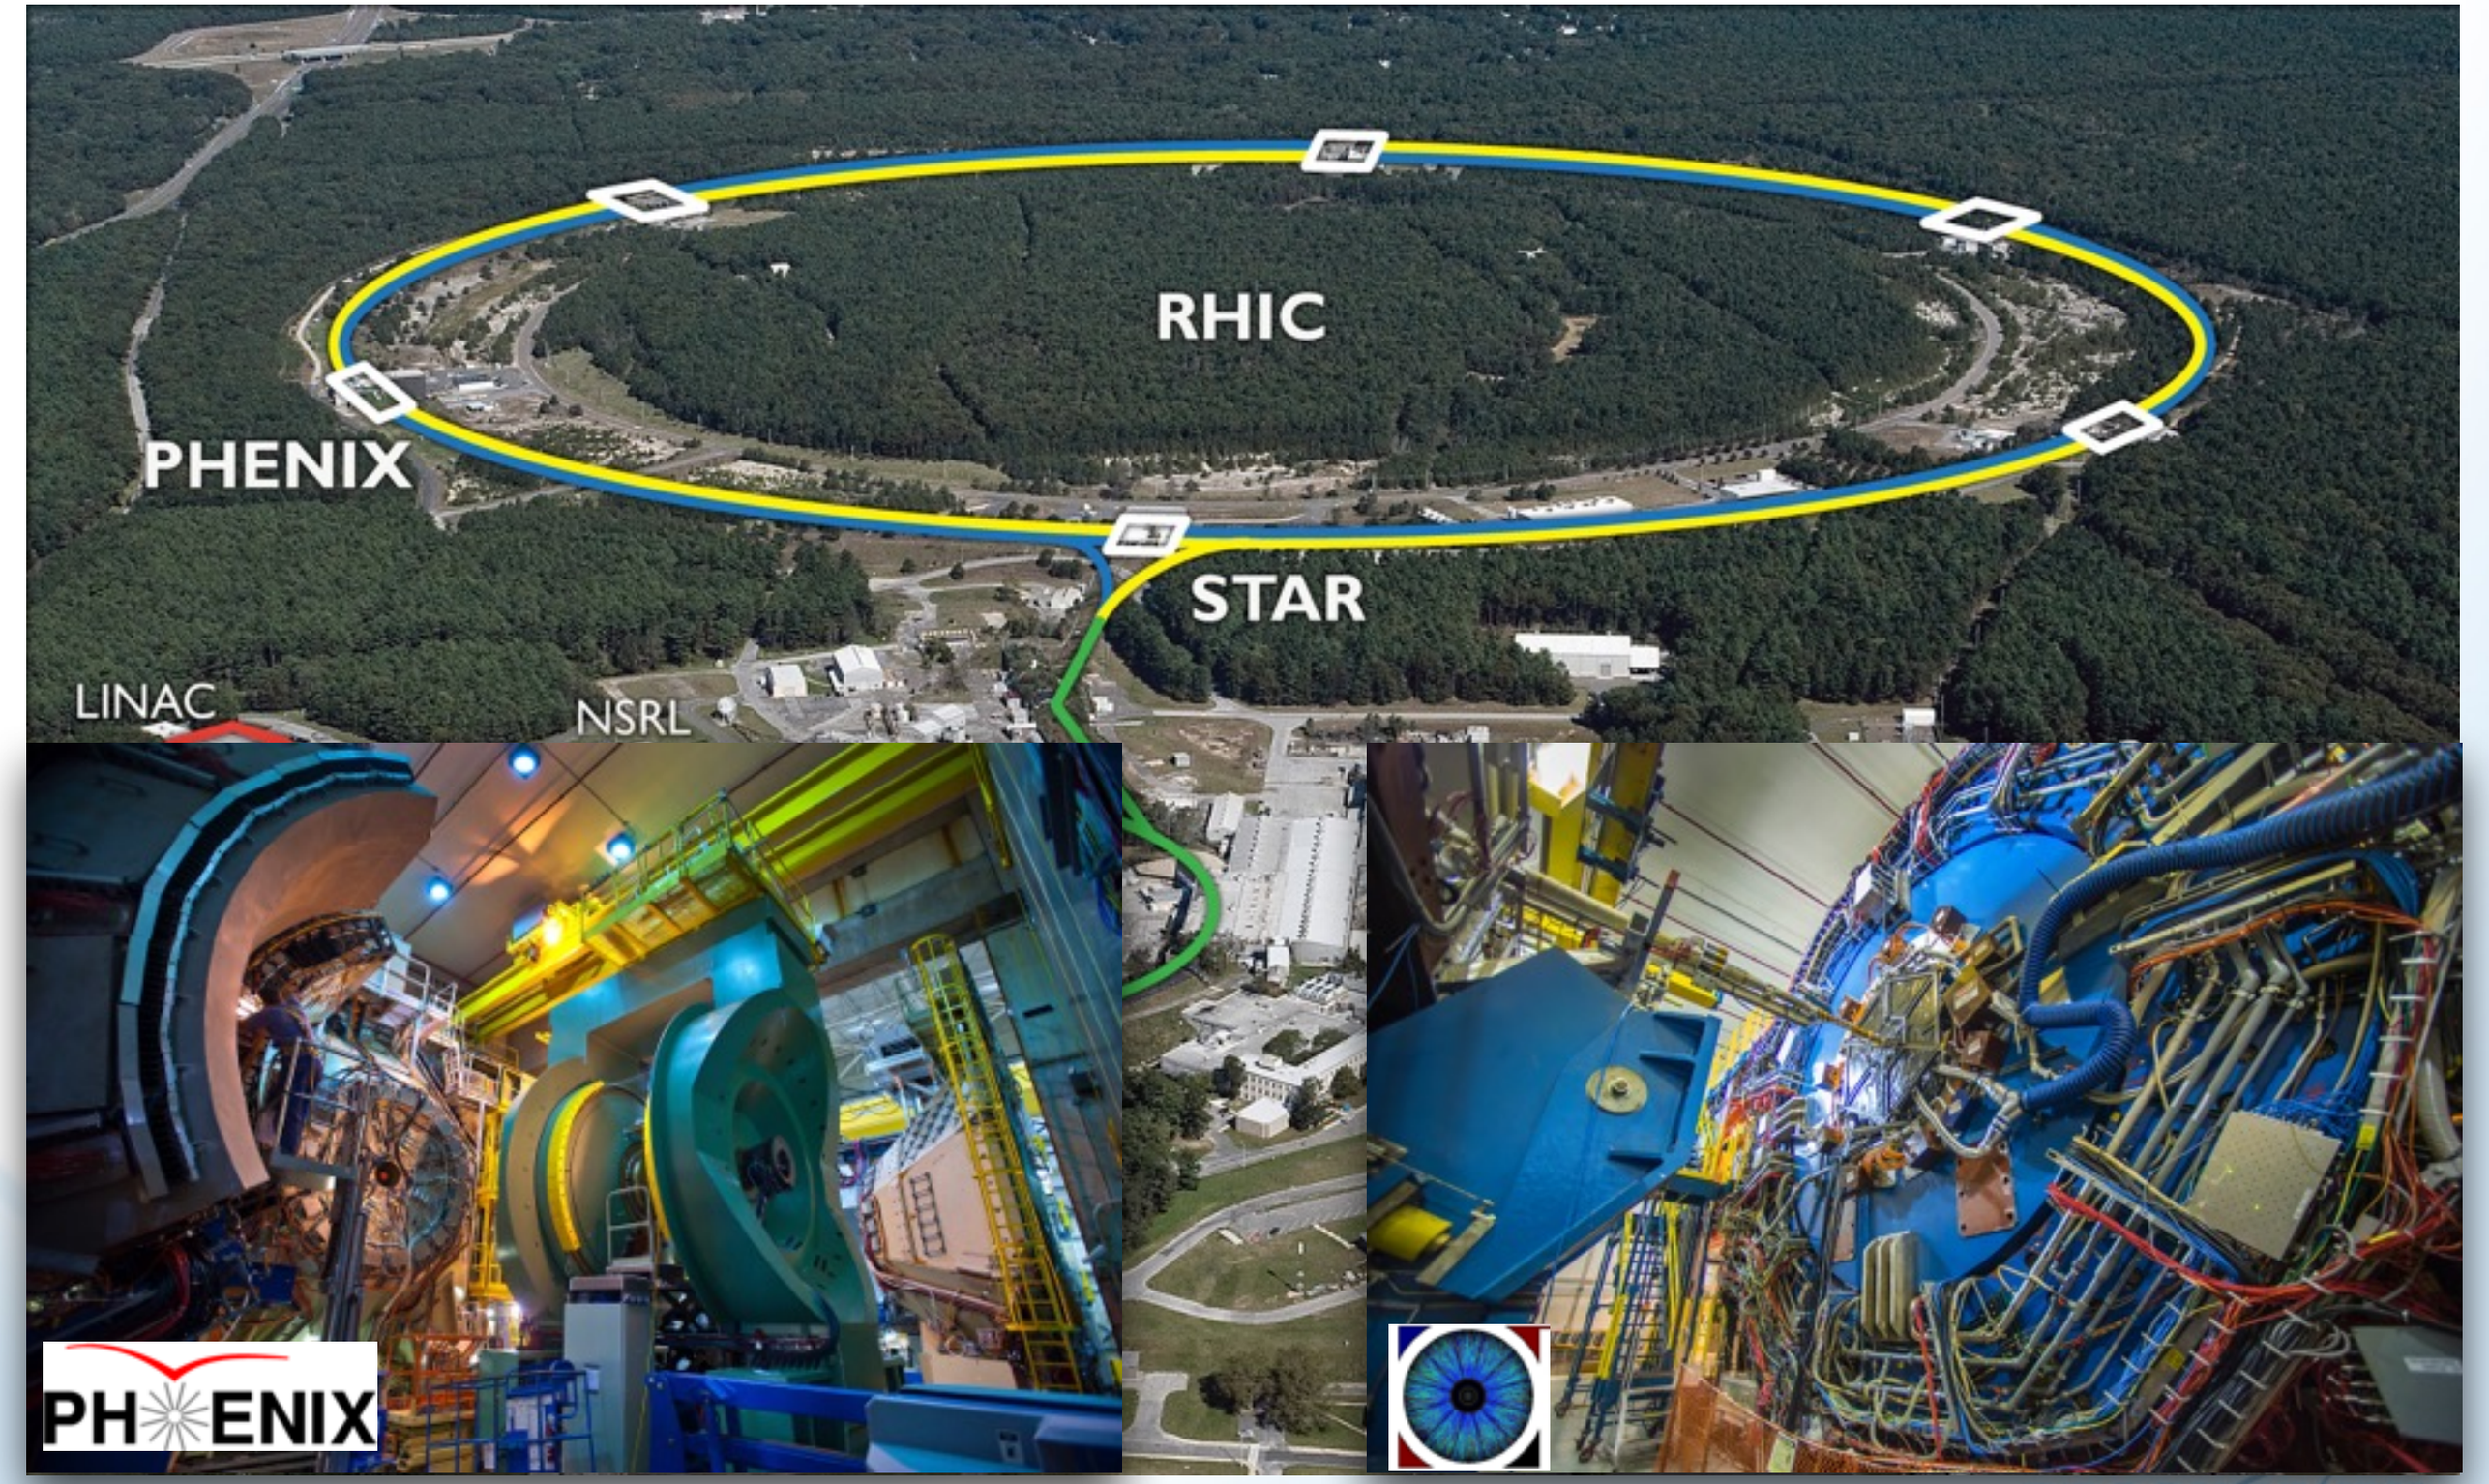

Scientific Computing at BNL - I

- **RACF : The RHIC and ATLAS Computing Facility,**
  - The historical core of scientific computing
  - A world class facility for High Energy Physics and Nuclear Physics
  - Grew organically to support computational needs of Nuclear Physics, QCD (theory), RHIC and US-ATLAS communities at large
  - The RHIC Tier 0
  - The US ATLAS Tier 1
  - Computing support for other experiments/projects of the Nuclear & physics department

Some of the current and future experiments of the physics department

DEEP UNDERGROUND NEUTRINO EXPERIMENT

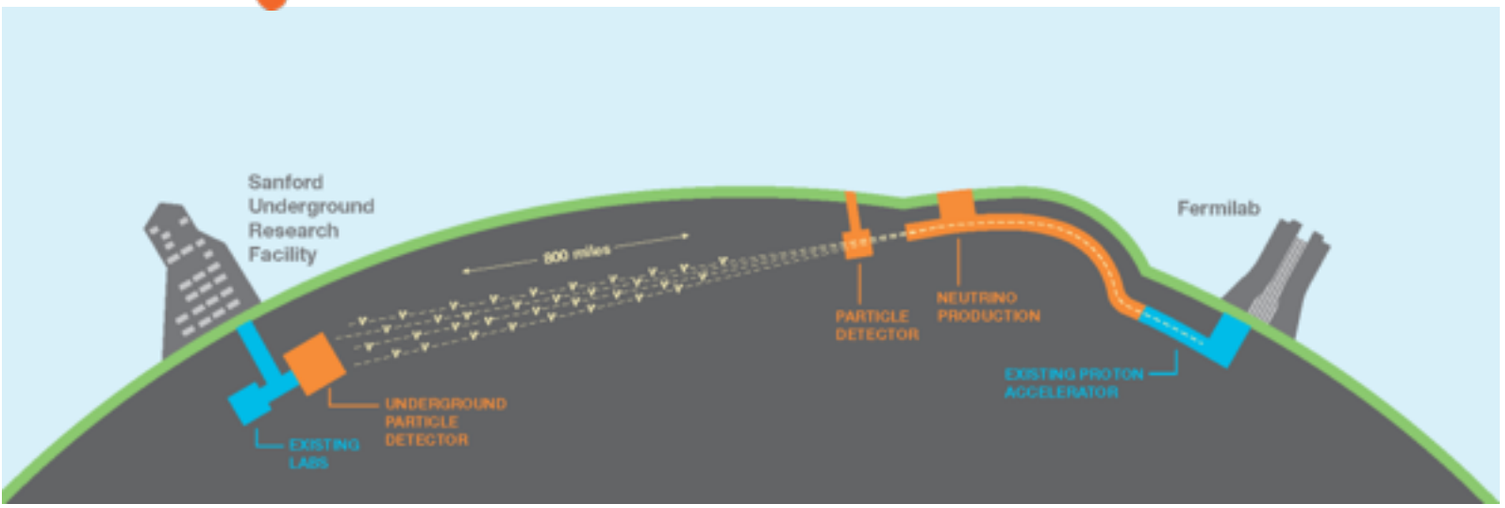

Dark Energy Survey

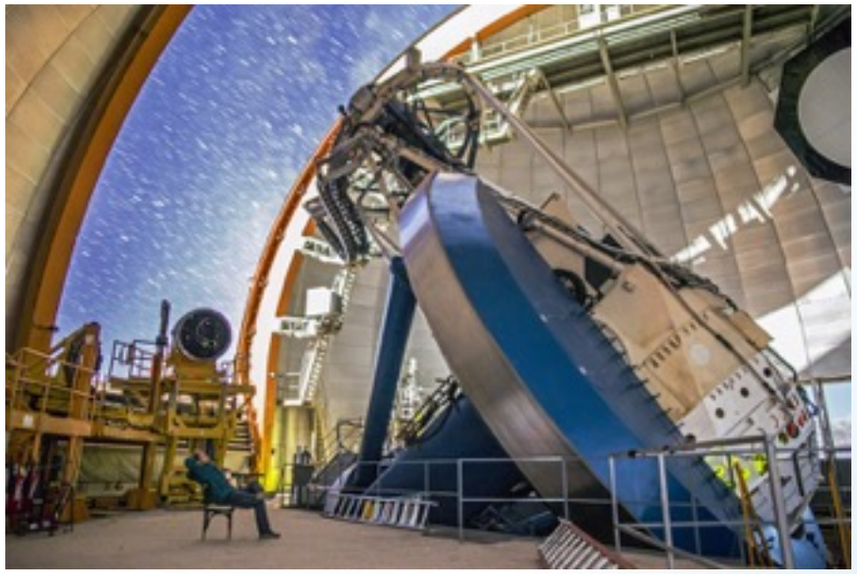

LSST

Large Synoptic Survey Telescope (LSST)

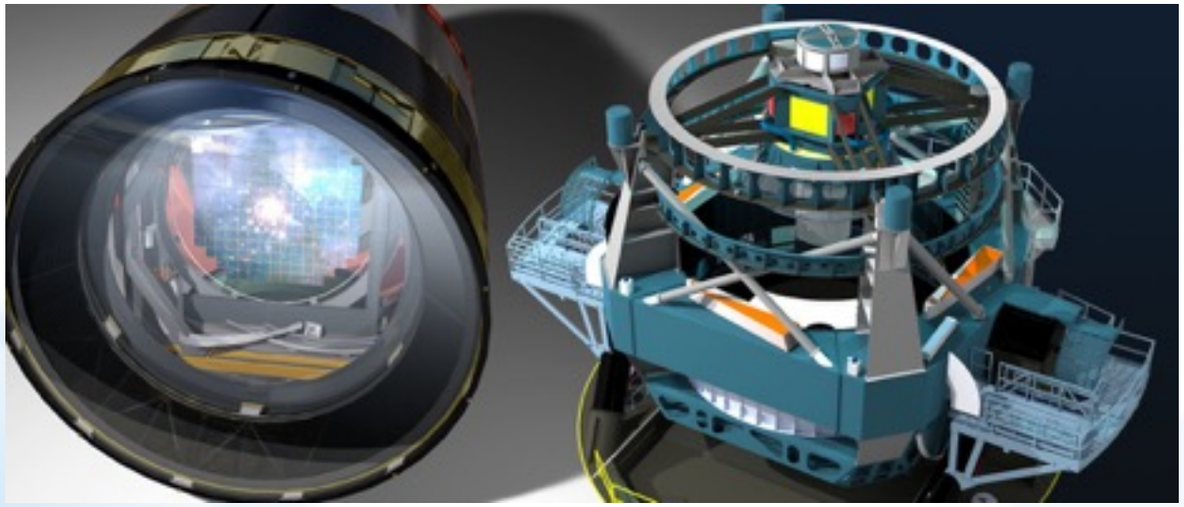

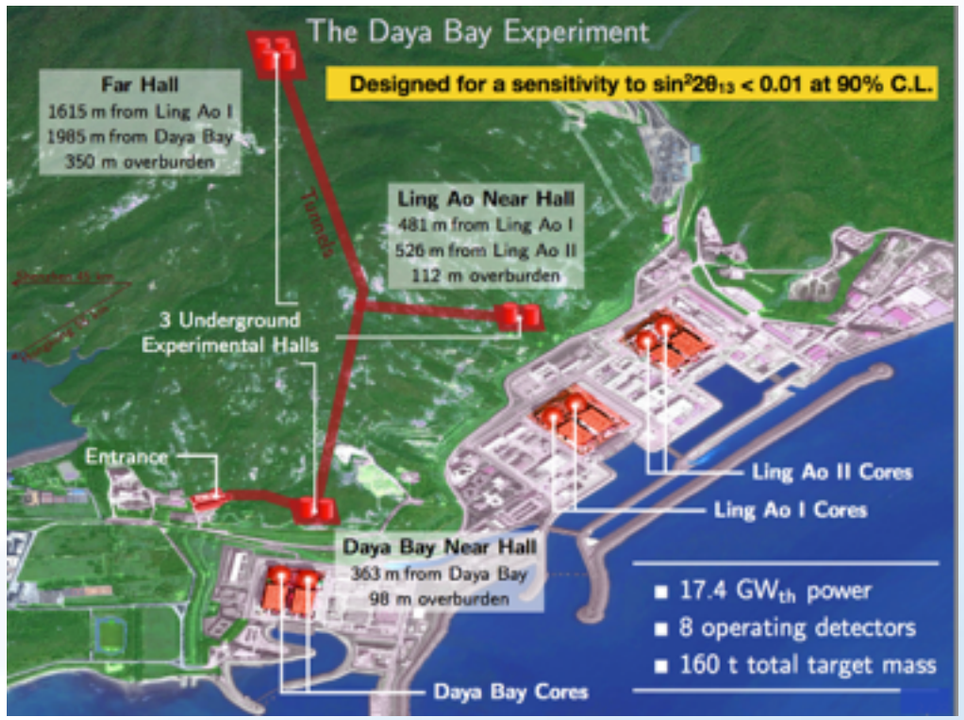

Some of the facilities on BNL site

Scientific Computing at BNL - II

- **CSI : BNL's Computational Science Initiative**
  - Established in December 2015 to provide an integrating umbrella for all computer science and applied mathematics research and services.
  - Provide a focal point for leading data science research
  - Educate the next generation of expert data scientists
  - Translate data science research advances into tools and expertise that lead to measurable scientific progress
  - Provide a State of the Art Data, Computing and Network Infrastructure to meet research and operational requirements for BNL and its Partners
- The computing center of CSI: **SDCC** (Scientific Data and Computing Center) is housed in RACF and operated by RACF staff

Computational Science Initiative : CSI

- Leverage laboratory investments in scientific computing across multiple programs
- Partners : universities (Columbia, Cornell, New York University, Stony Brook, and Yale) and companies including IBM Research.
- **SDCC:** BNL Scientific Data and Computing Center includes components from
  - Laboratory's Institutional Cluster
  - CFN (Center for Functional Nano-materials)
  - Atmospheric Radiation Measurement
  - USQCD
  - …

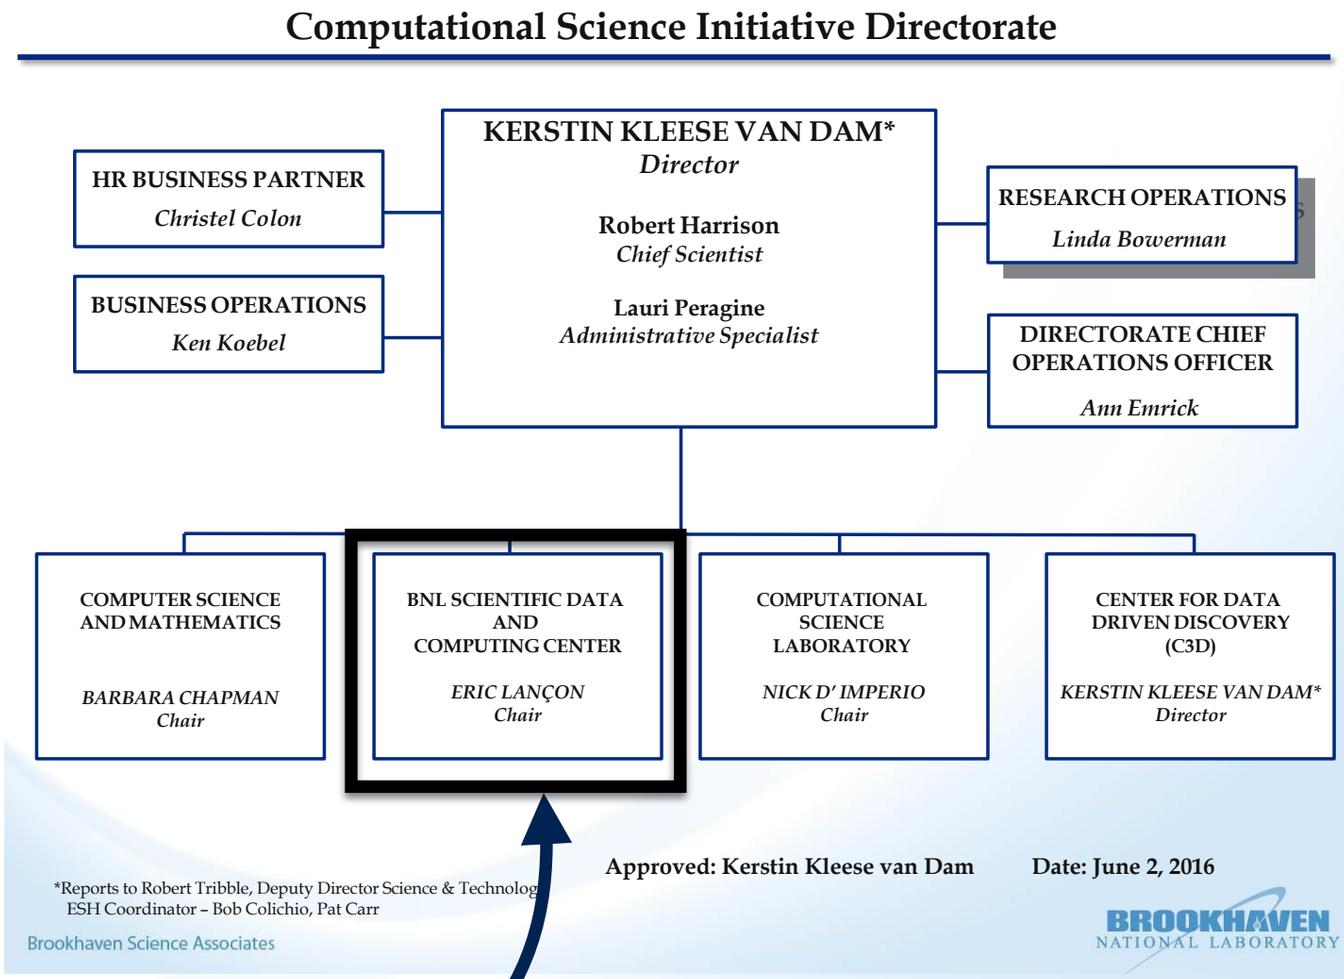

SDCC

SDCC today

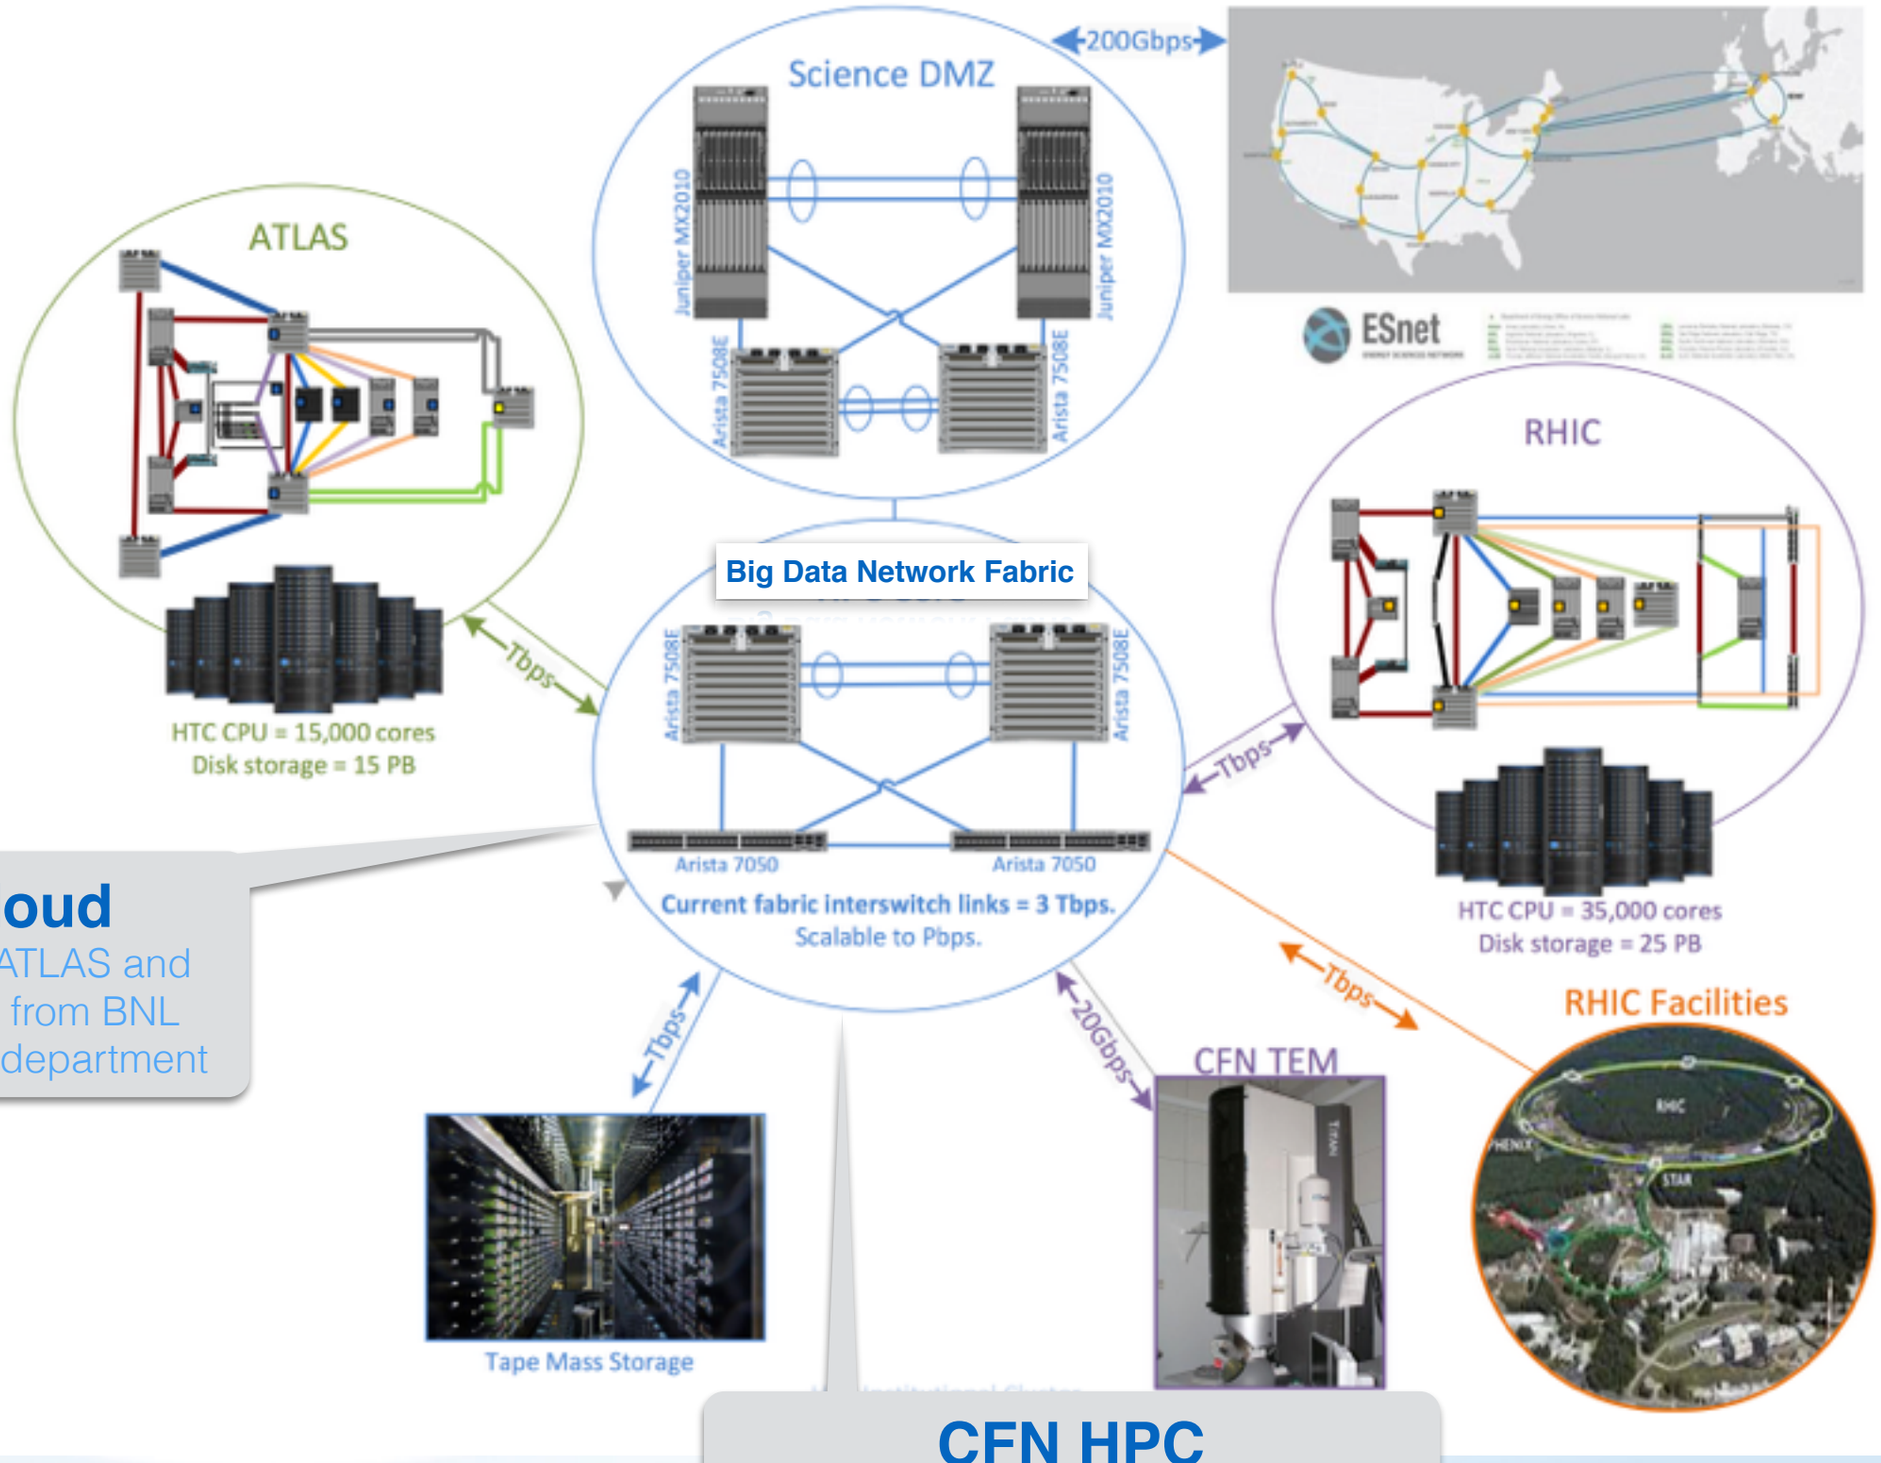

BNL Cloud
Used by OSG, ATLAS and various groups from BNL outside physics department

CFN HPC
(Center for Functional Nano-materials)

SDCC in Numbers

- **70k CPU cores (~100k in 2017)**
  - HPC 22% of capacity

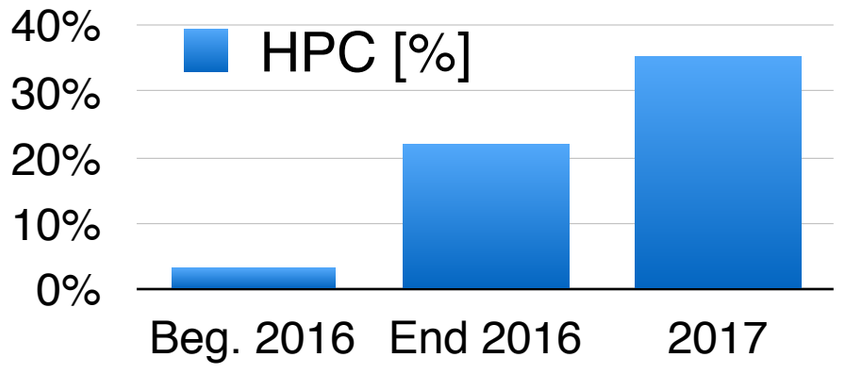

- **~50 PB of disk storage**
  - of various technologies (dCache, Xrootd, Ceph, Gpfs,…)

- **~80 PB of tape storage**
  - 4th HPSS (High Performance Storage System) site worldwide
  - first within the US(1)

(1) http://www.hpss-collaboration.org/learn_who_petabyte_data.shtml

70k cores

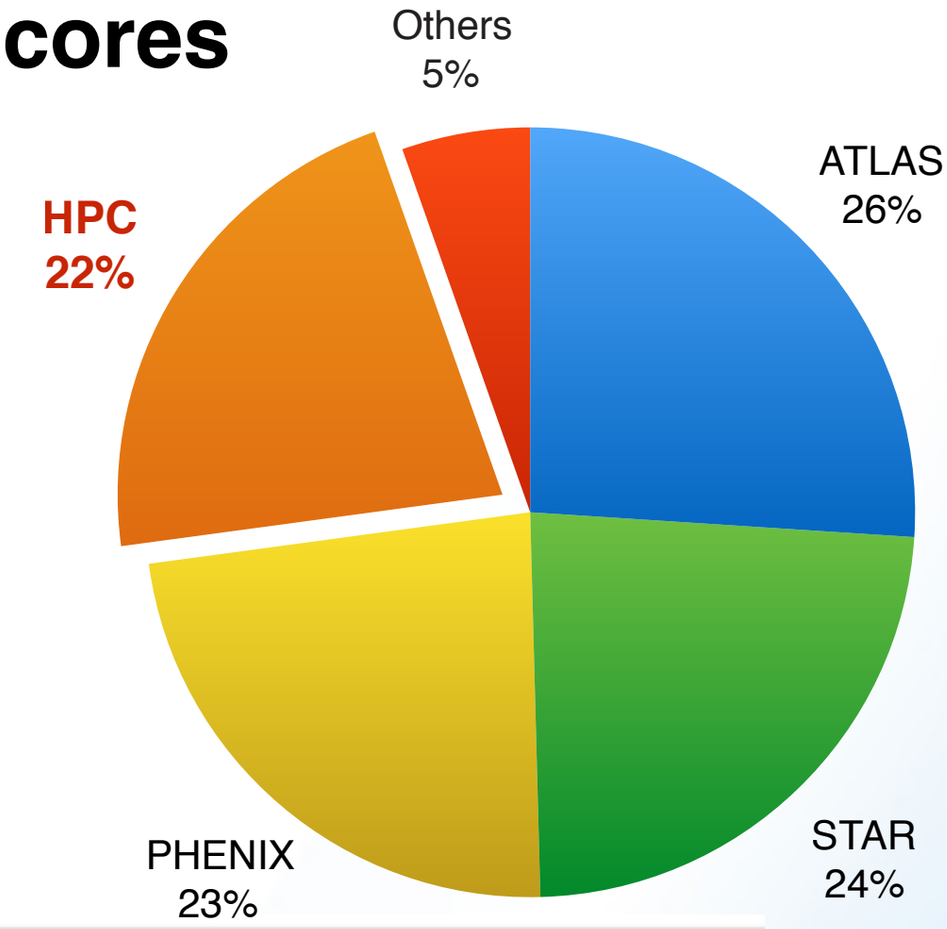

HPSS sites

Site
(ECMWF) European Centre for Medium-Range Weather Forecasts
(NOAA-RD) National Oceanic and Atmospheric Administration Research & Development
(UKMO) United Kingdom Met Office
(BNL) Brookhaven National Laboratory
(LBNL-User) Lawrence Berkley National Laboratory - User
(LANL-Secure) Los Alamos National Laboratory - Secure
(ORNL) Oak Ridge National Laboratory
(NCAR) National Center for Atmospheric Research

Expertise

- **HTC** : High Throughput Computing
  - Leadership in data driven computing
  - Supporting 1000s of users

- **Data storage and archiving**
  - Disk storage (dCache, XRoot, GPFS, Ceph, …)
  - Tape archive

- **Networking**
  - WAN : High performance Wide Area Network transfers (in collaboration with ESnet)
  - LAN : Local Area Network with Terabit scale capability

- **Cloud and virtualisation technologies**
  - Pioneer in collaboration with AWS (Amazon Web Services) and Google
  - Object storage

The RHIC Tier 0

- Store and process data from RHIC experiments
- Provide analysis means for thousands of users
- Data preservation long after end of experiments

Performance in 2016

- Data transfer from experiments to facility

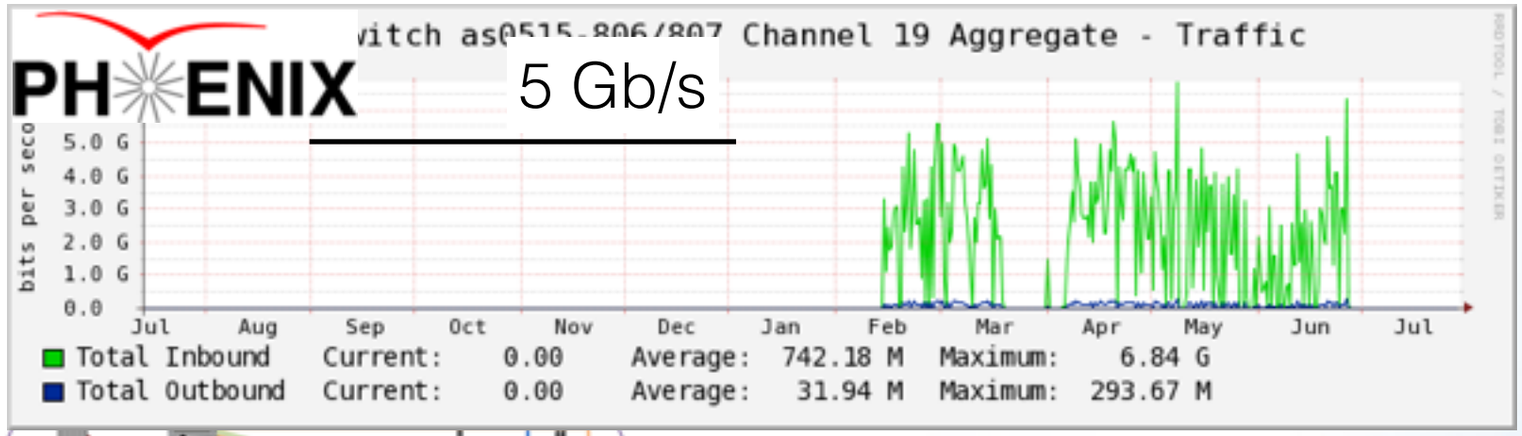

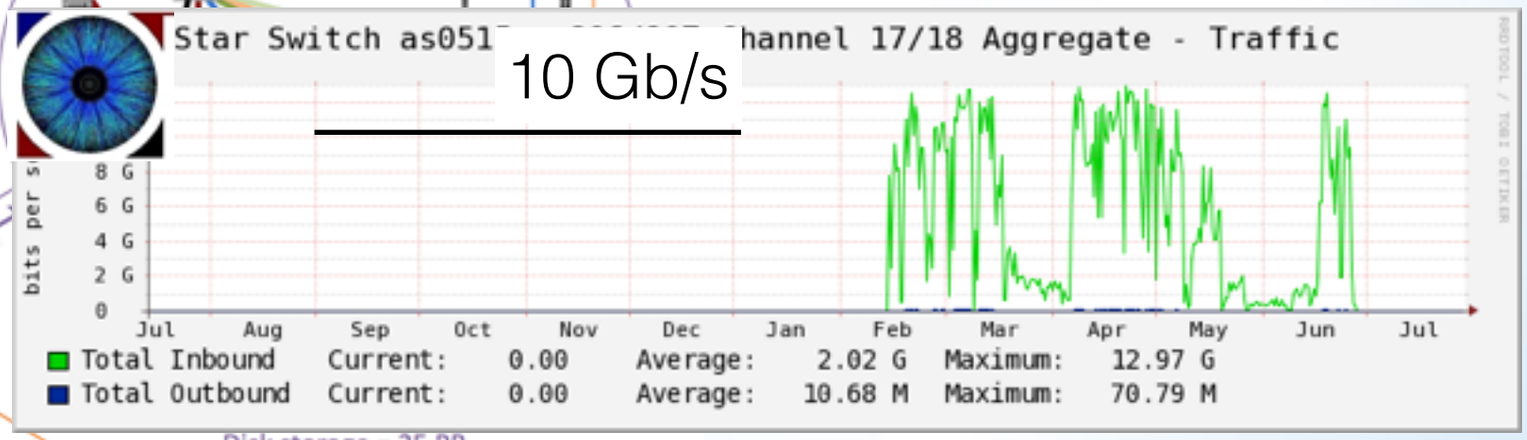

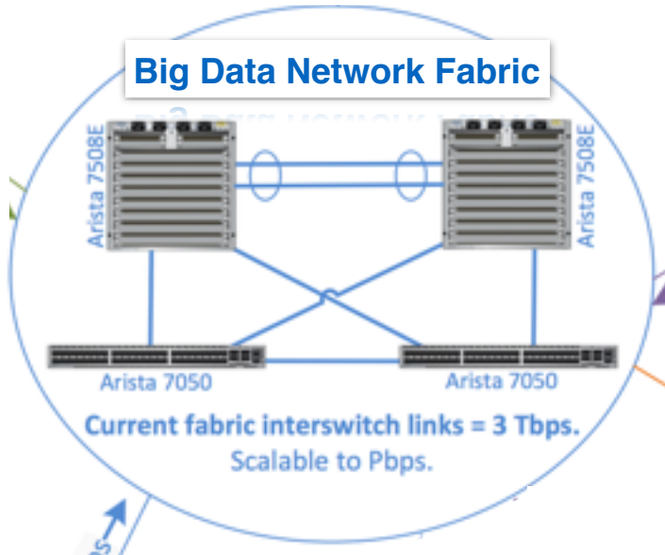

Performance in 2016

- Writing RAW data to tape

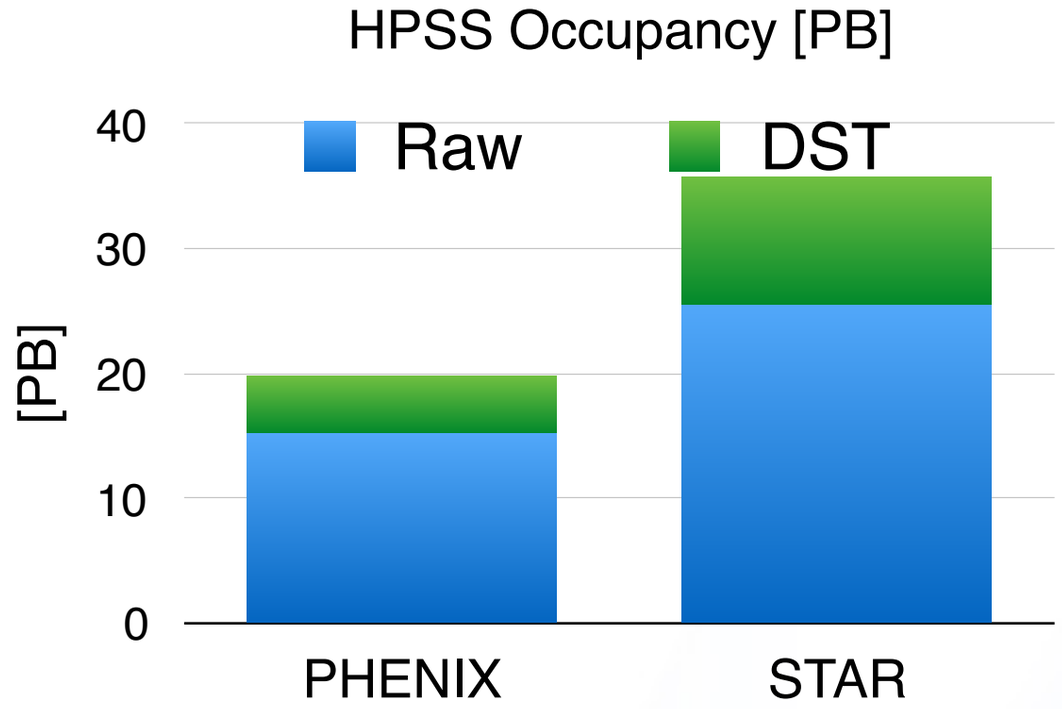

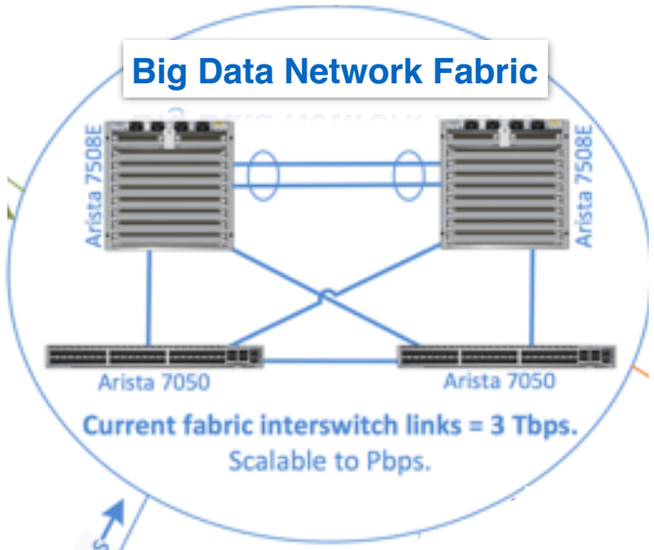

Tape Mass Storage

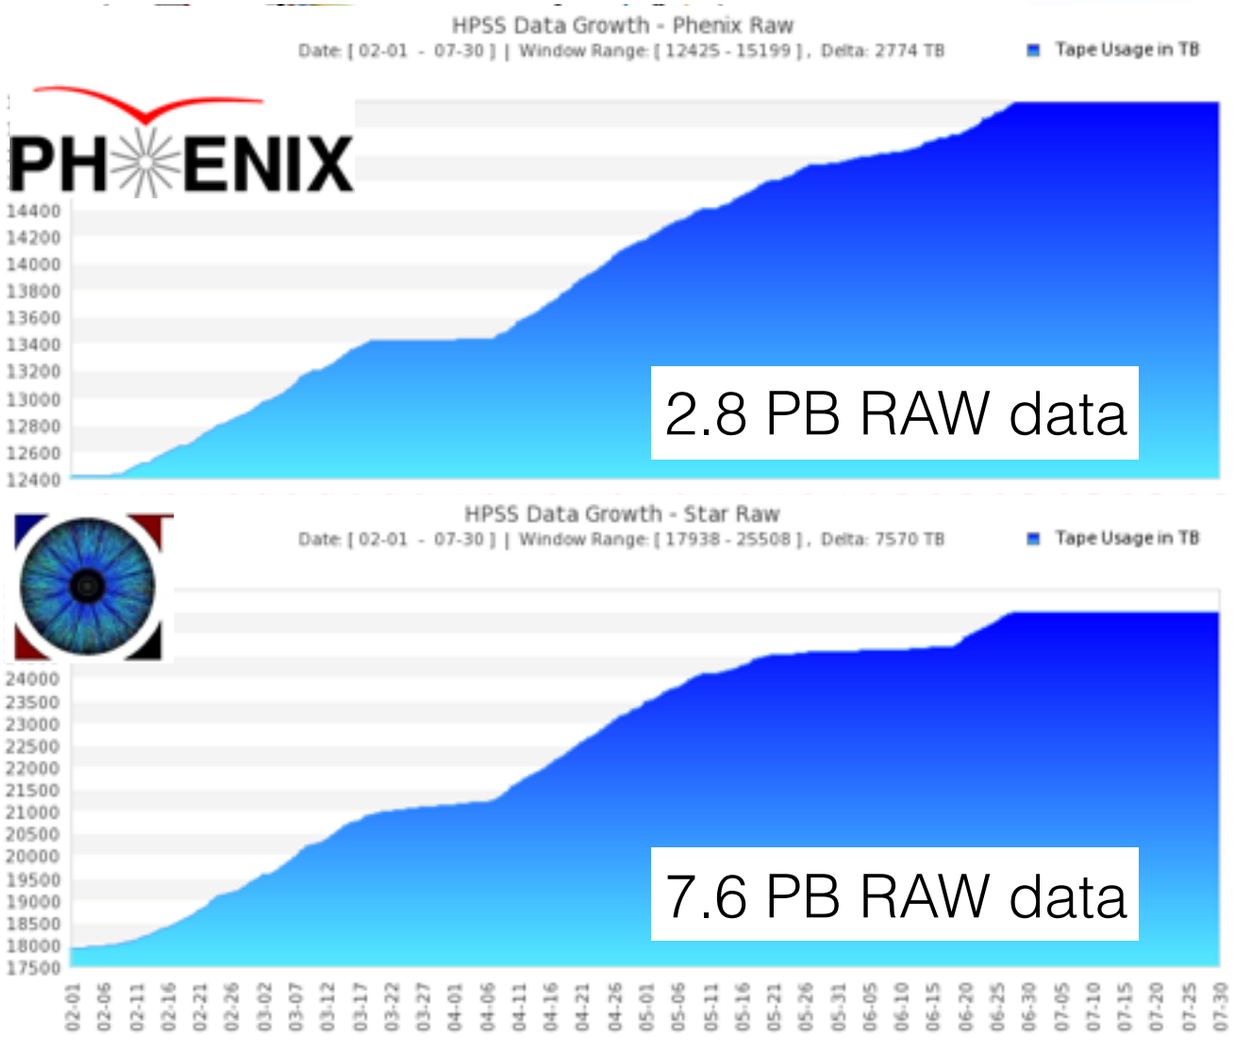

High Throughput Parallel Archiving

RHIC RUN 16 - STAR

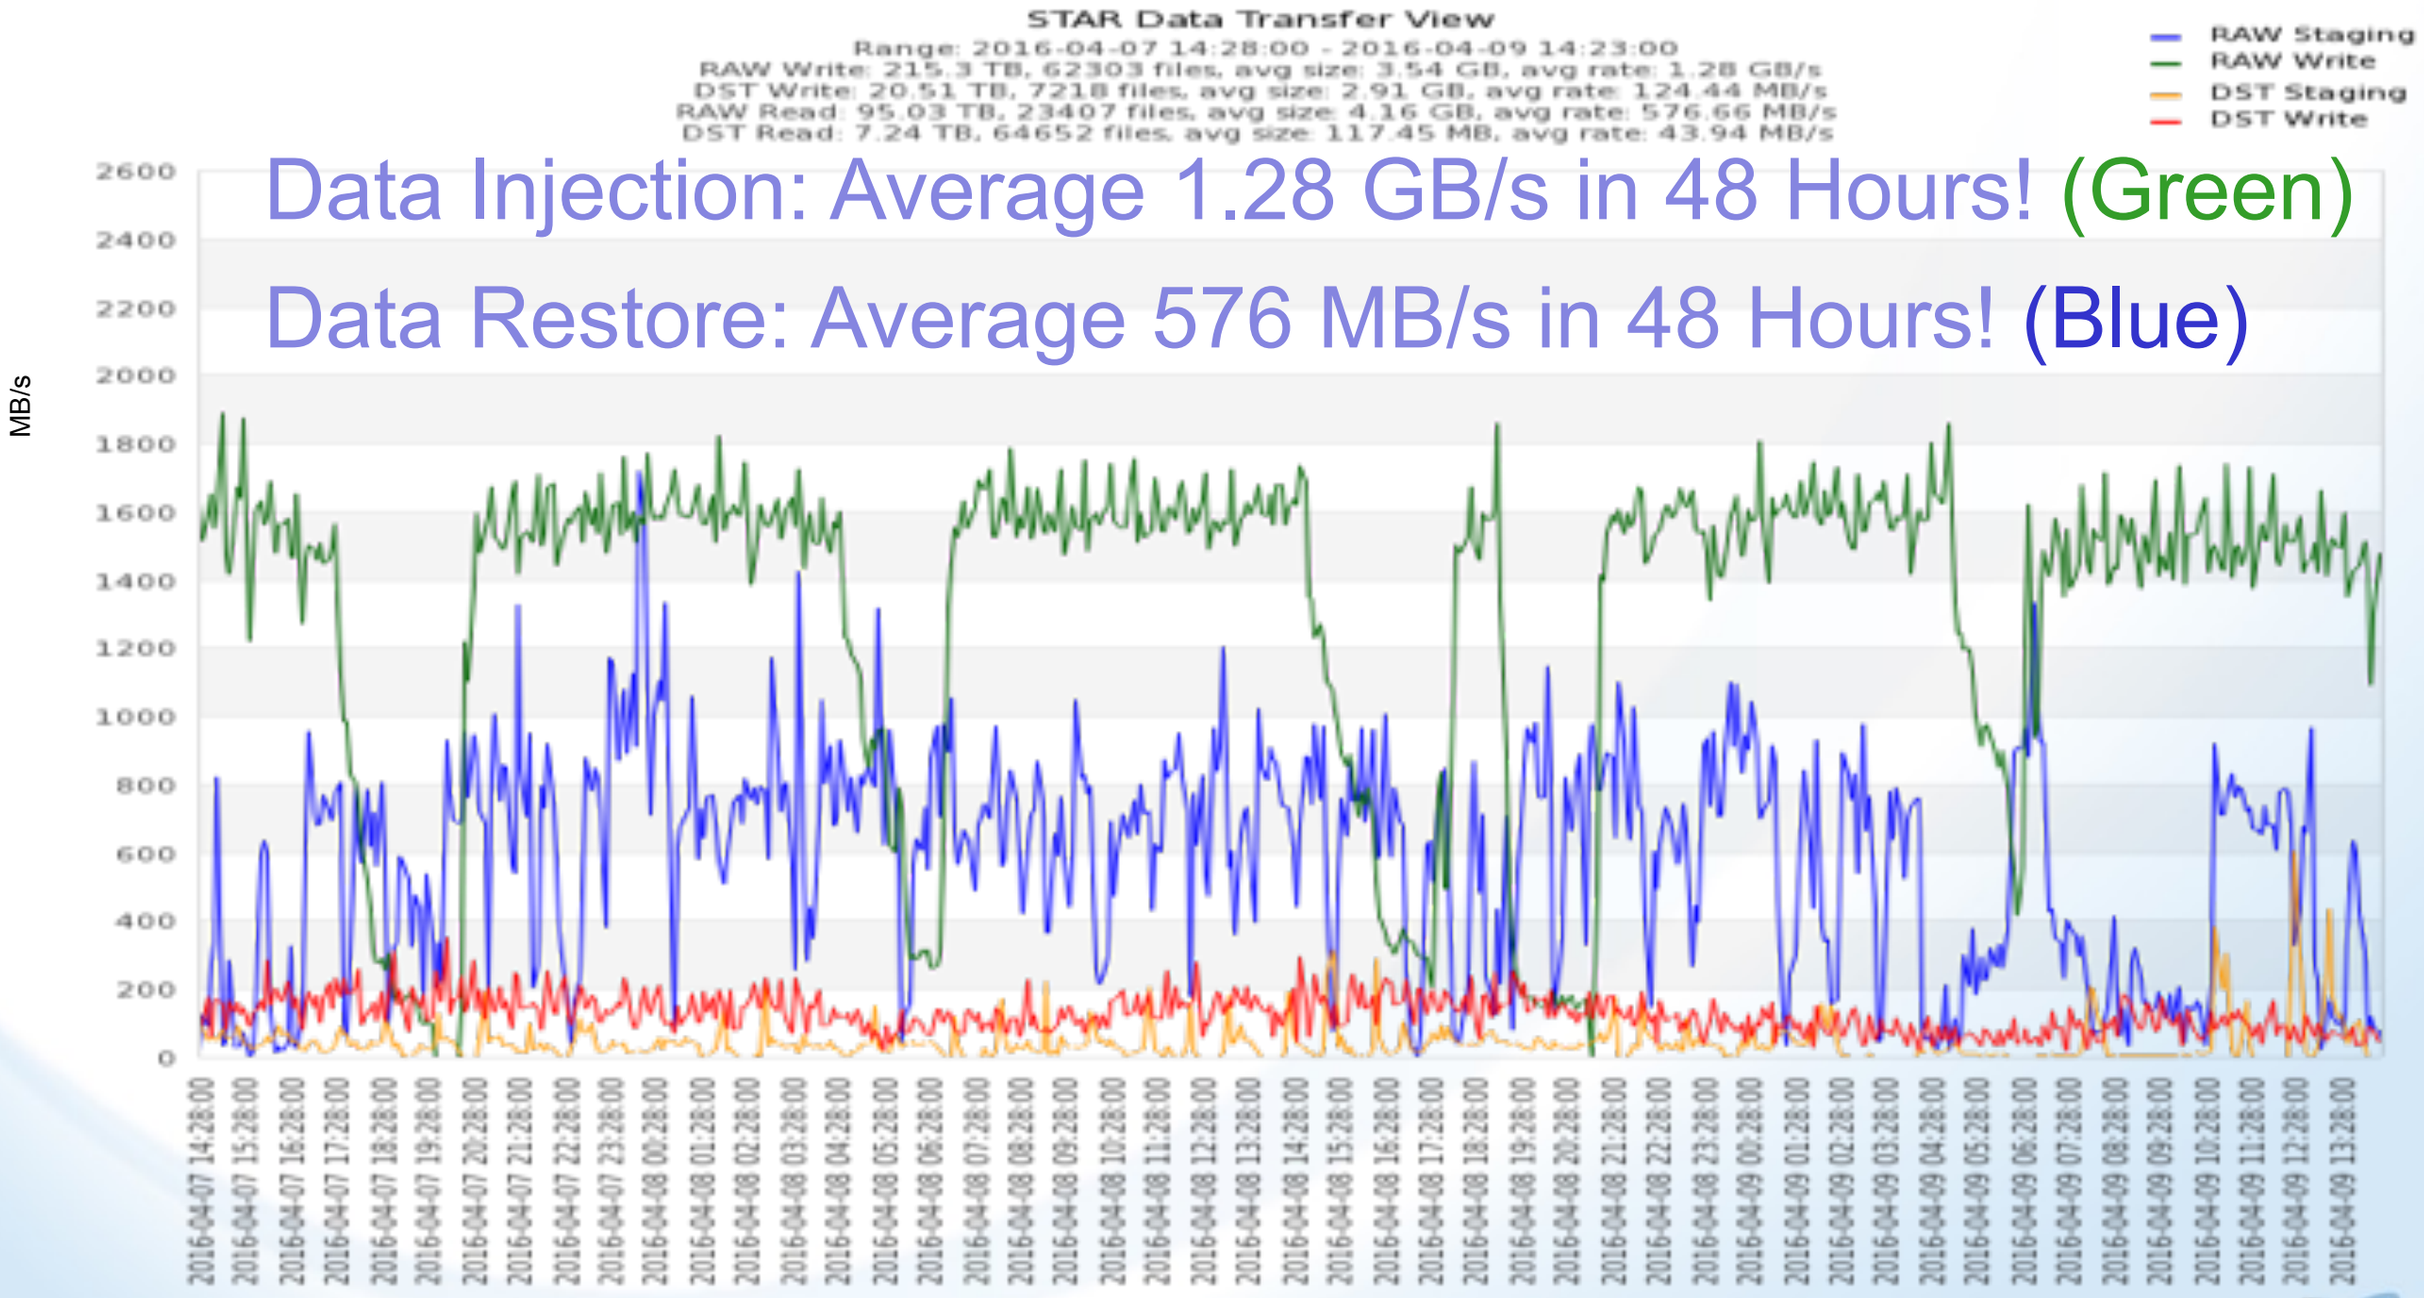

CERN Site

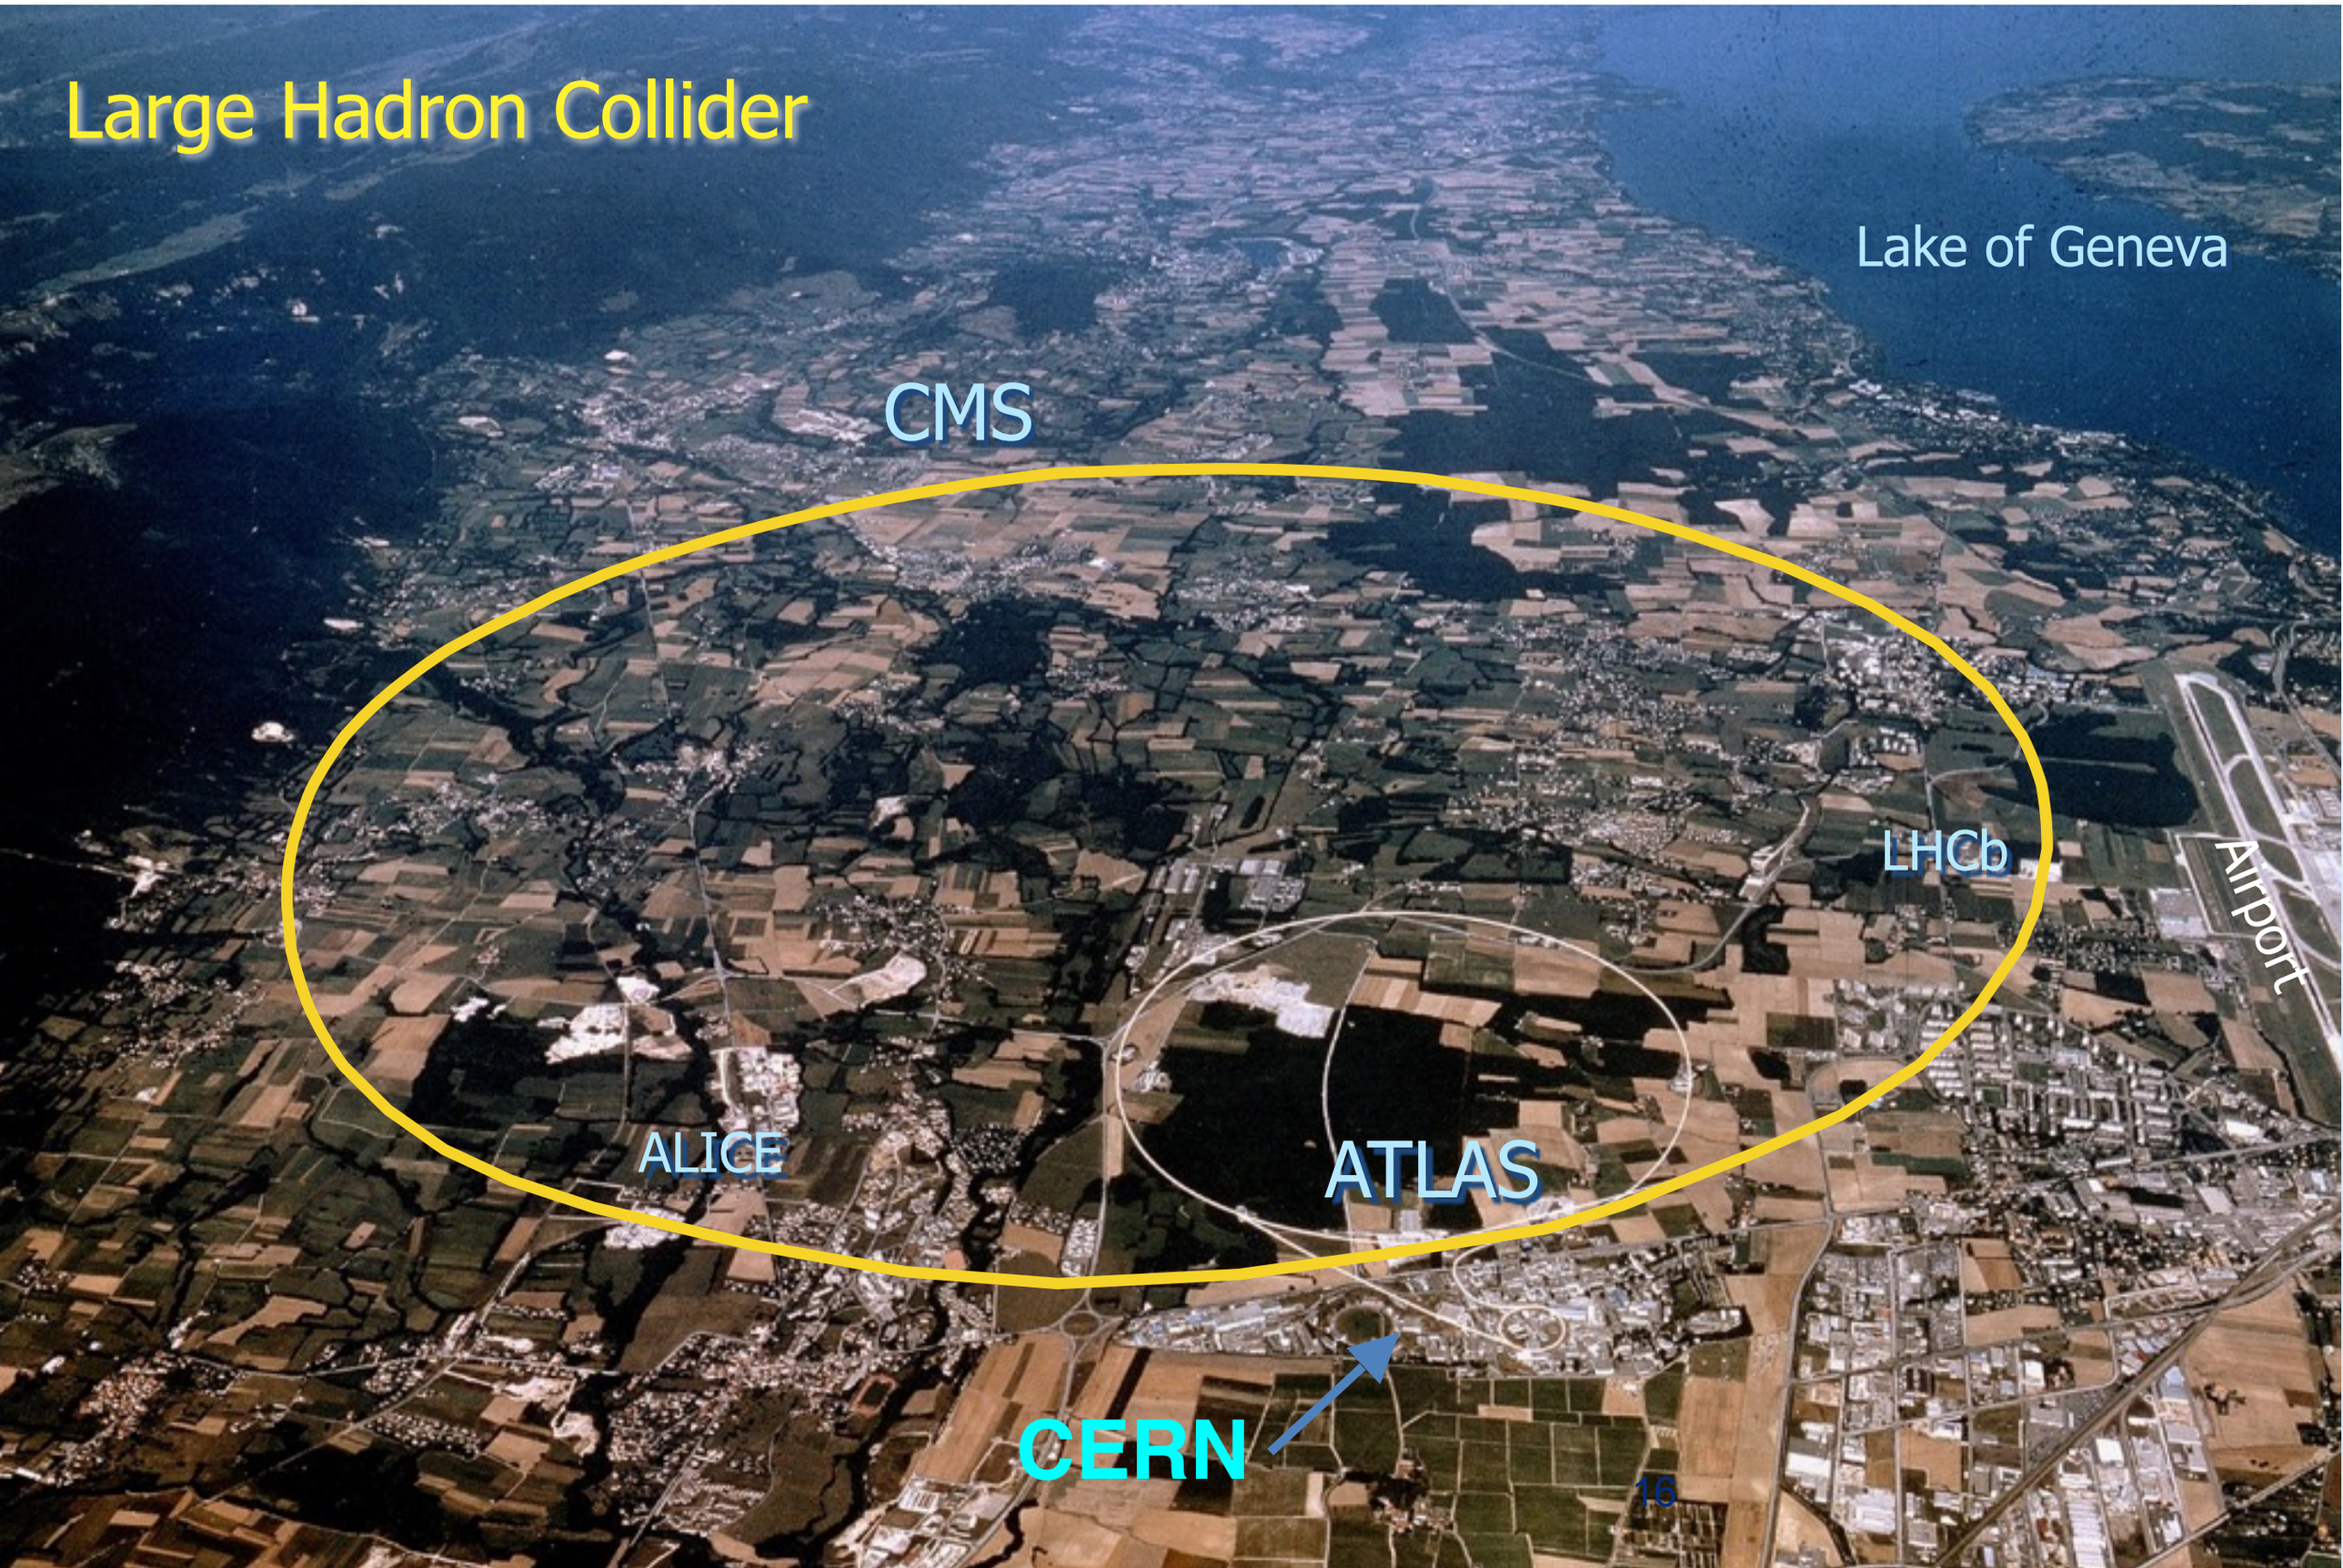
Large Hadron Collider
Lake of Geneva
CMS
LHCb
Airport
ALICE
ATLAS
CERN

This is ATLAS

Over 150 million sensors data

Register ~ 30 PB/y
at the pace of up to 4GB/s

ATLAS data are processed and analysed with ~250.000 CPU cores in ~80 centres in 30 countries

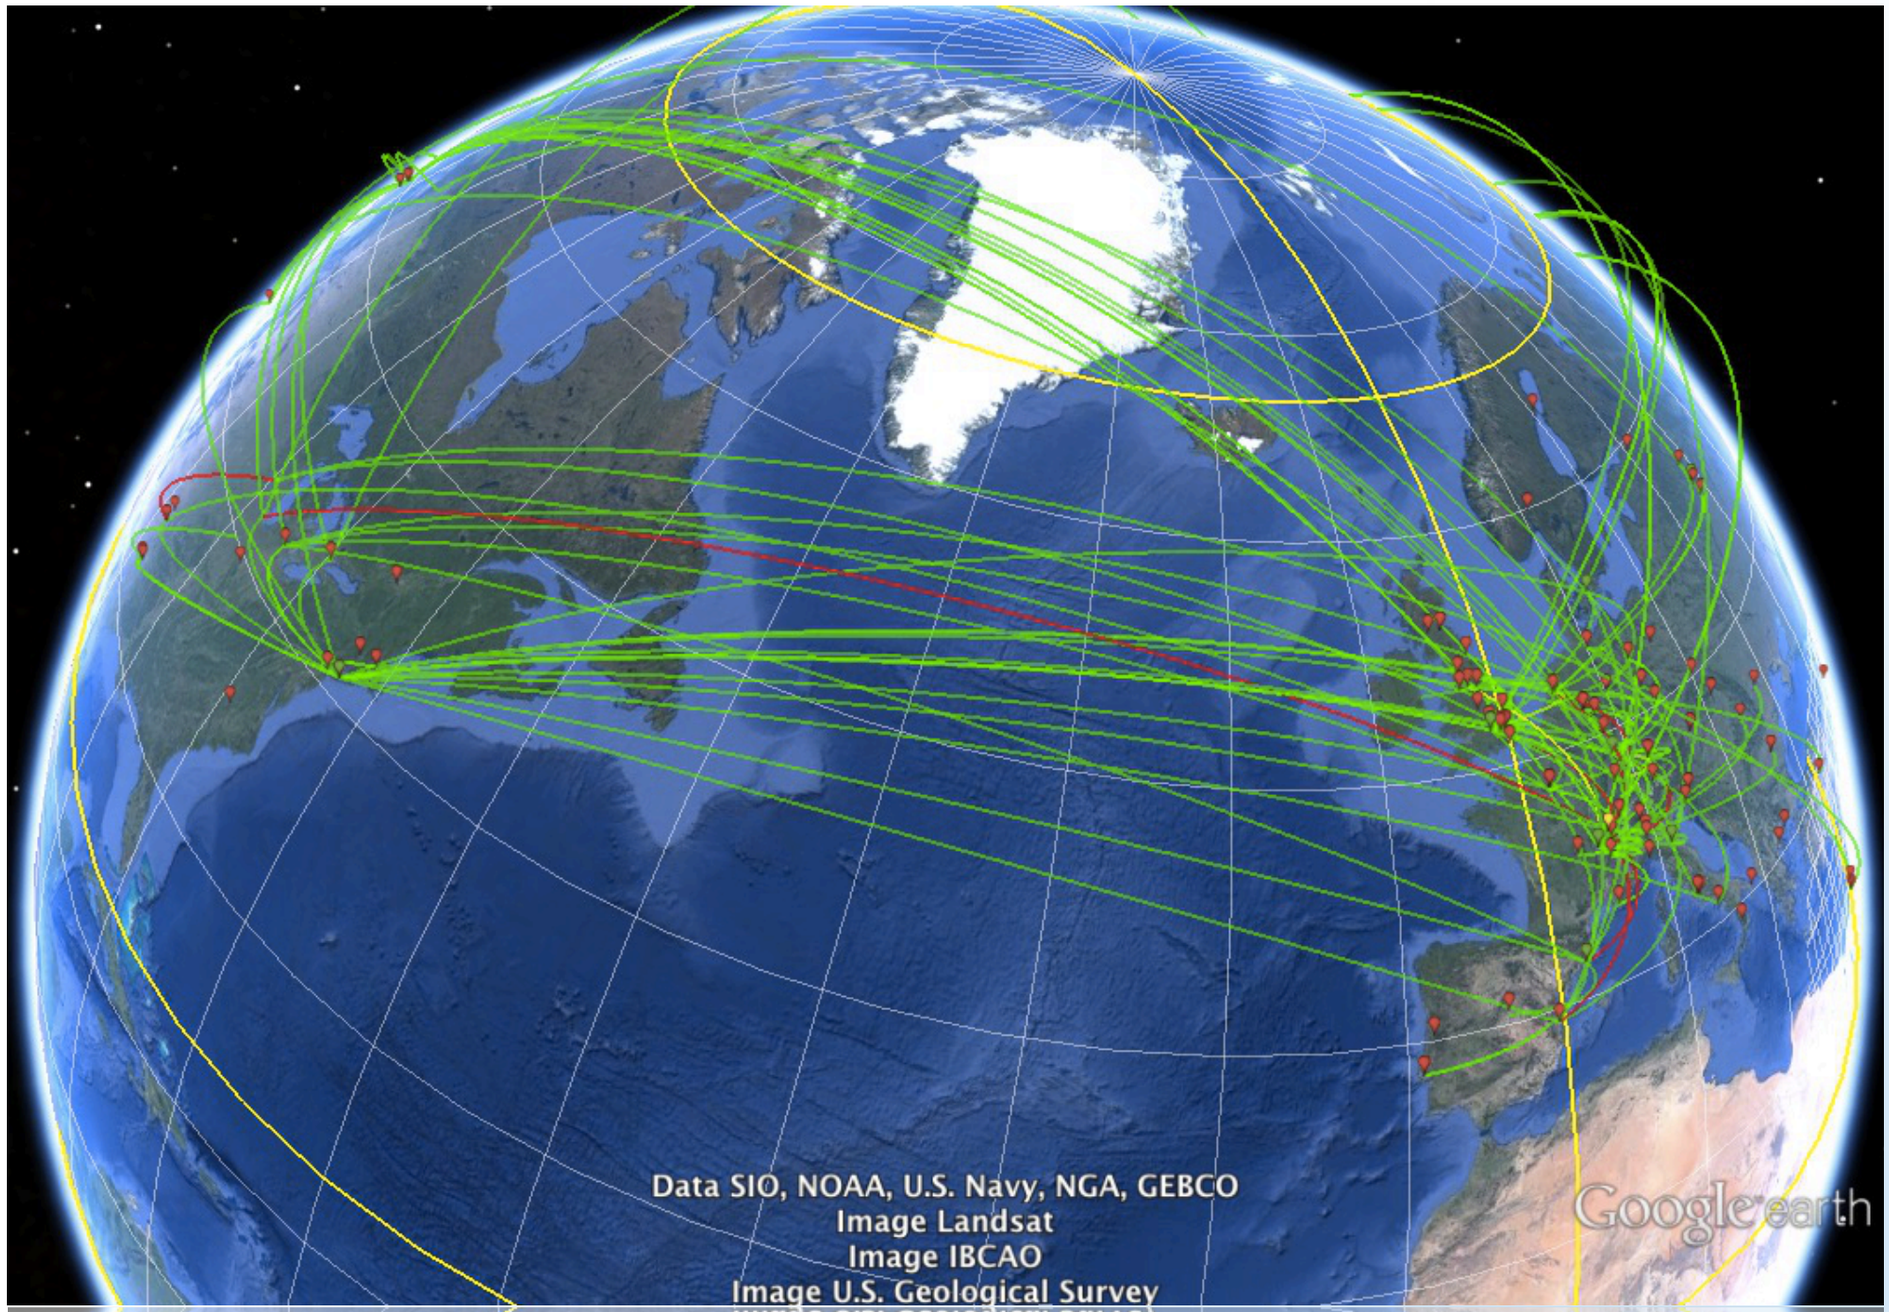

The US ATLAS Tier 1

- **99%** service availability requirement (MoU)
- **~25%** of ATLAS Tier 1 computing capacity worldwide
- Store RAW data from LHC and from simulation
- Distribute data to the 5 US Tier 2 sites
- Analysis center for US physicists

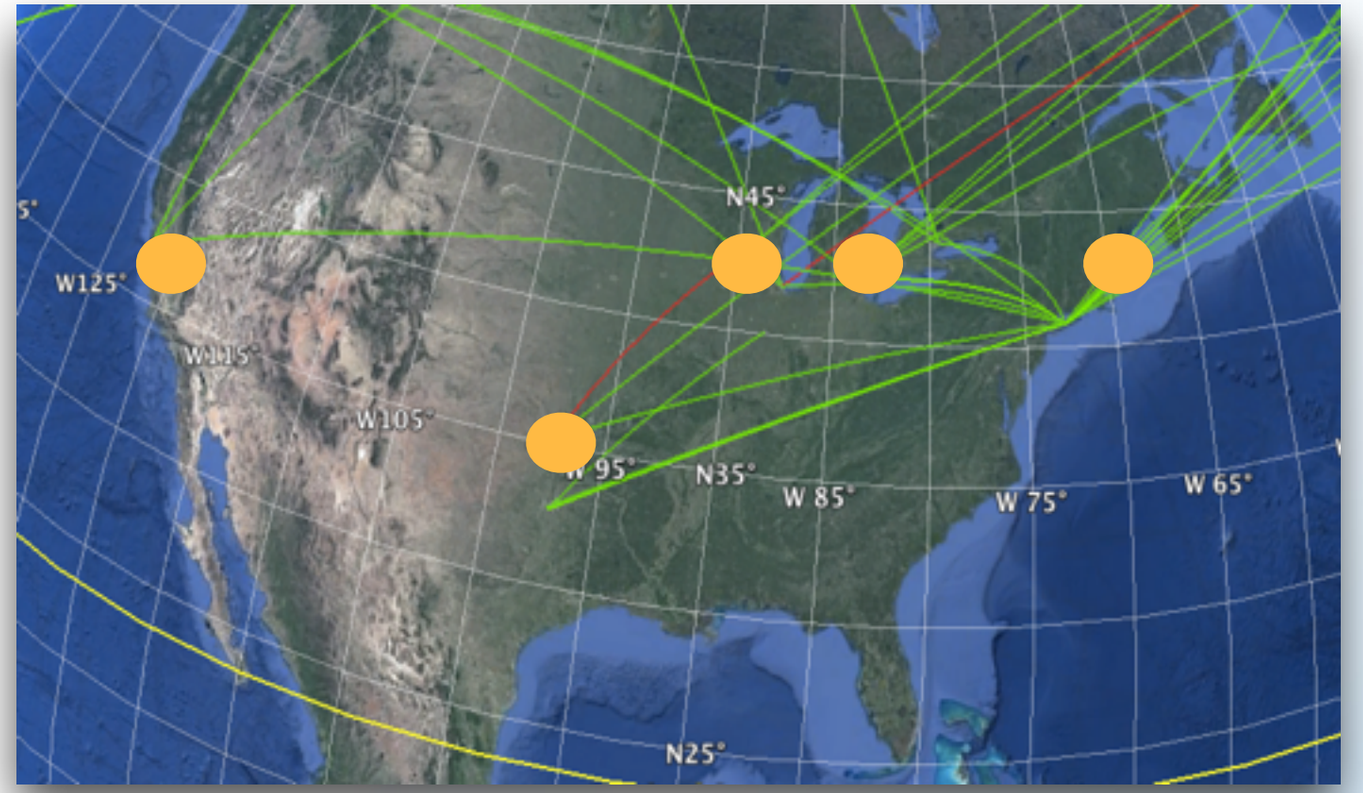

ATLAS: Big Data Science

- Expertise in storing, processing, distributing and providing resources for analysis of big data samples
- Over a year
  - Data import : 30 PB
  - Data export : 32 PB
  - 150 PB processed

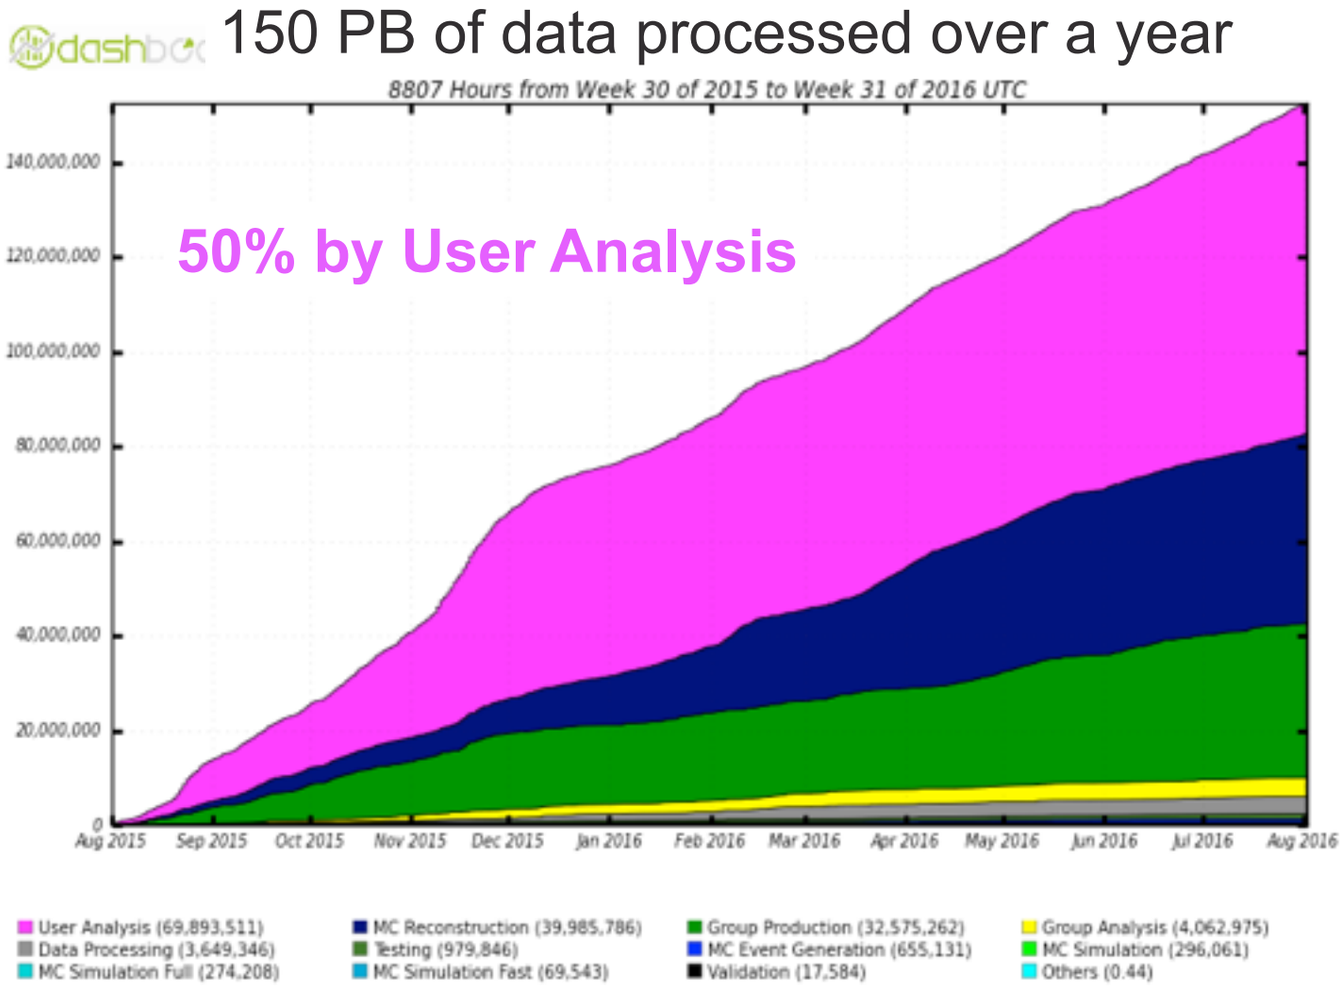

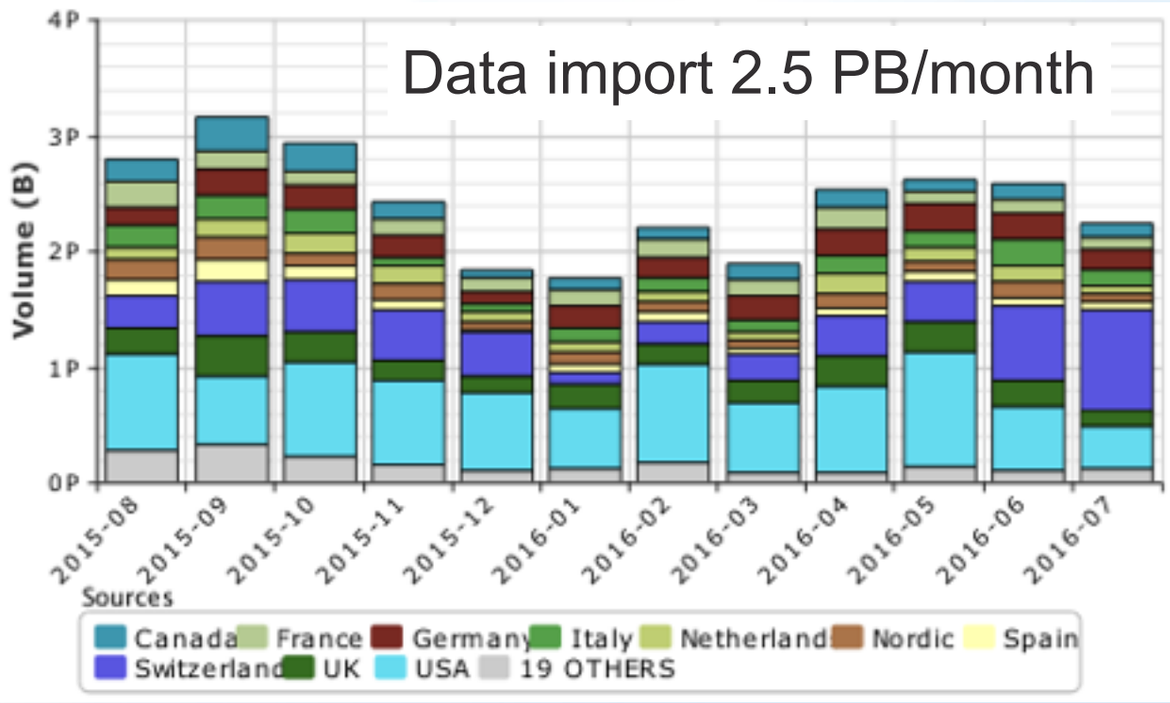

ATLAS: BNL and DOE HPCs

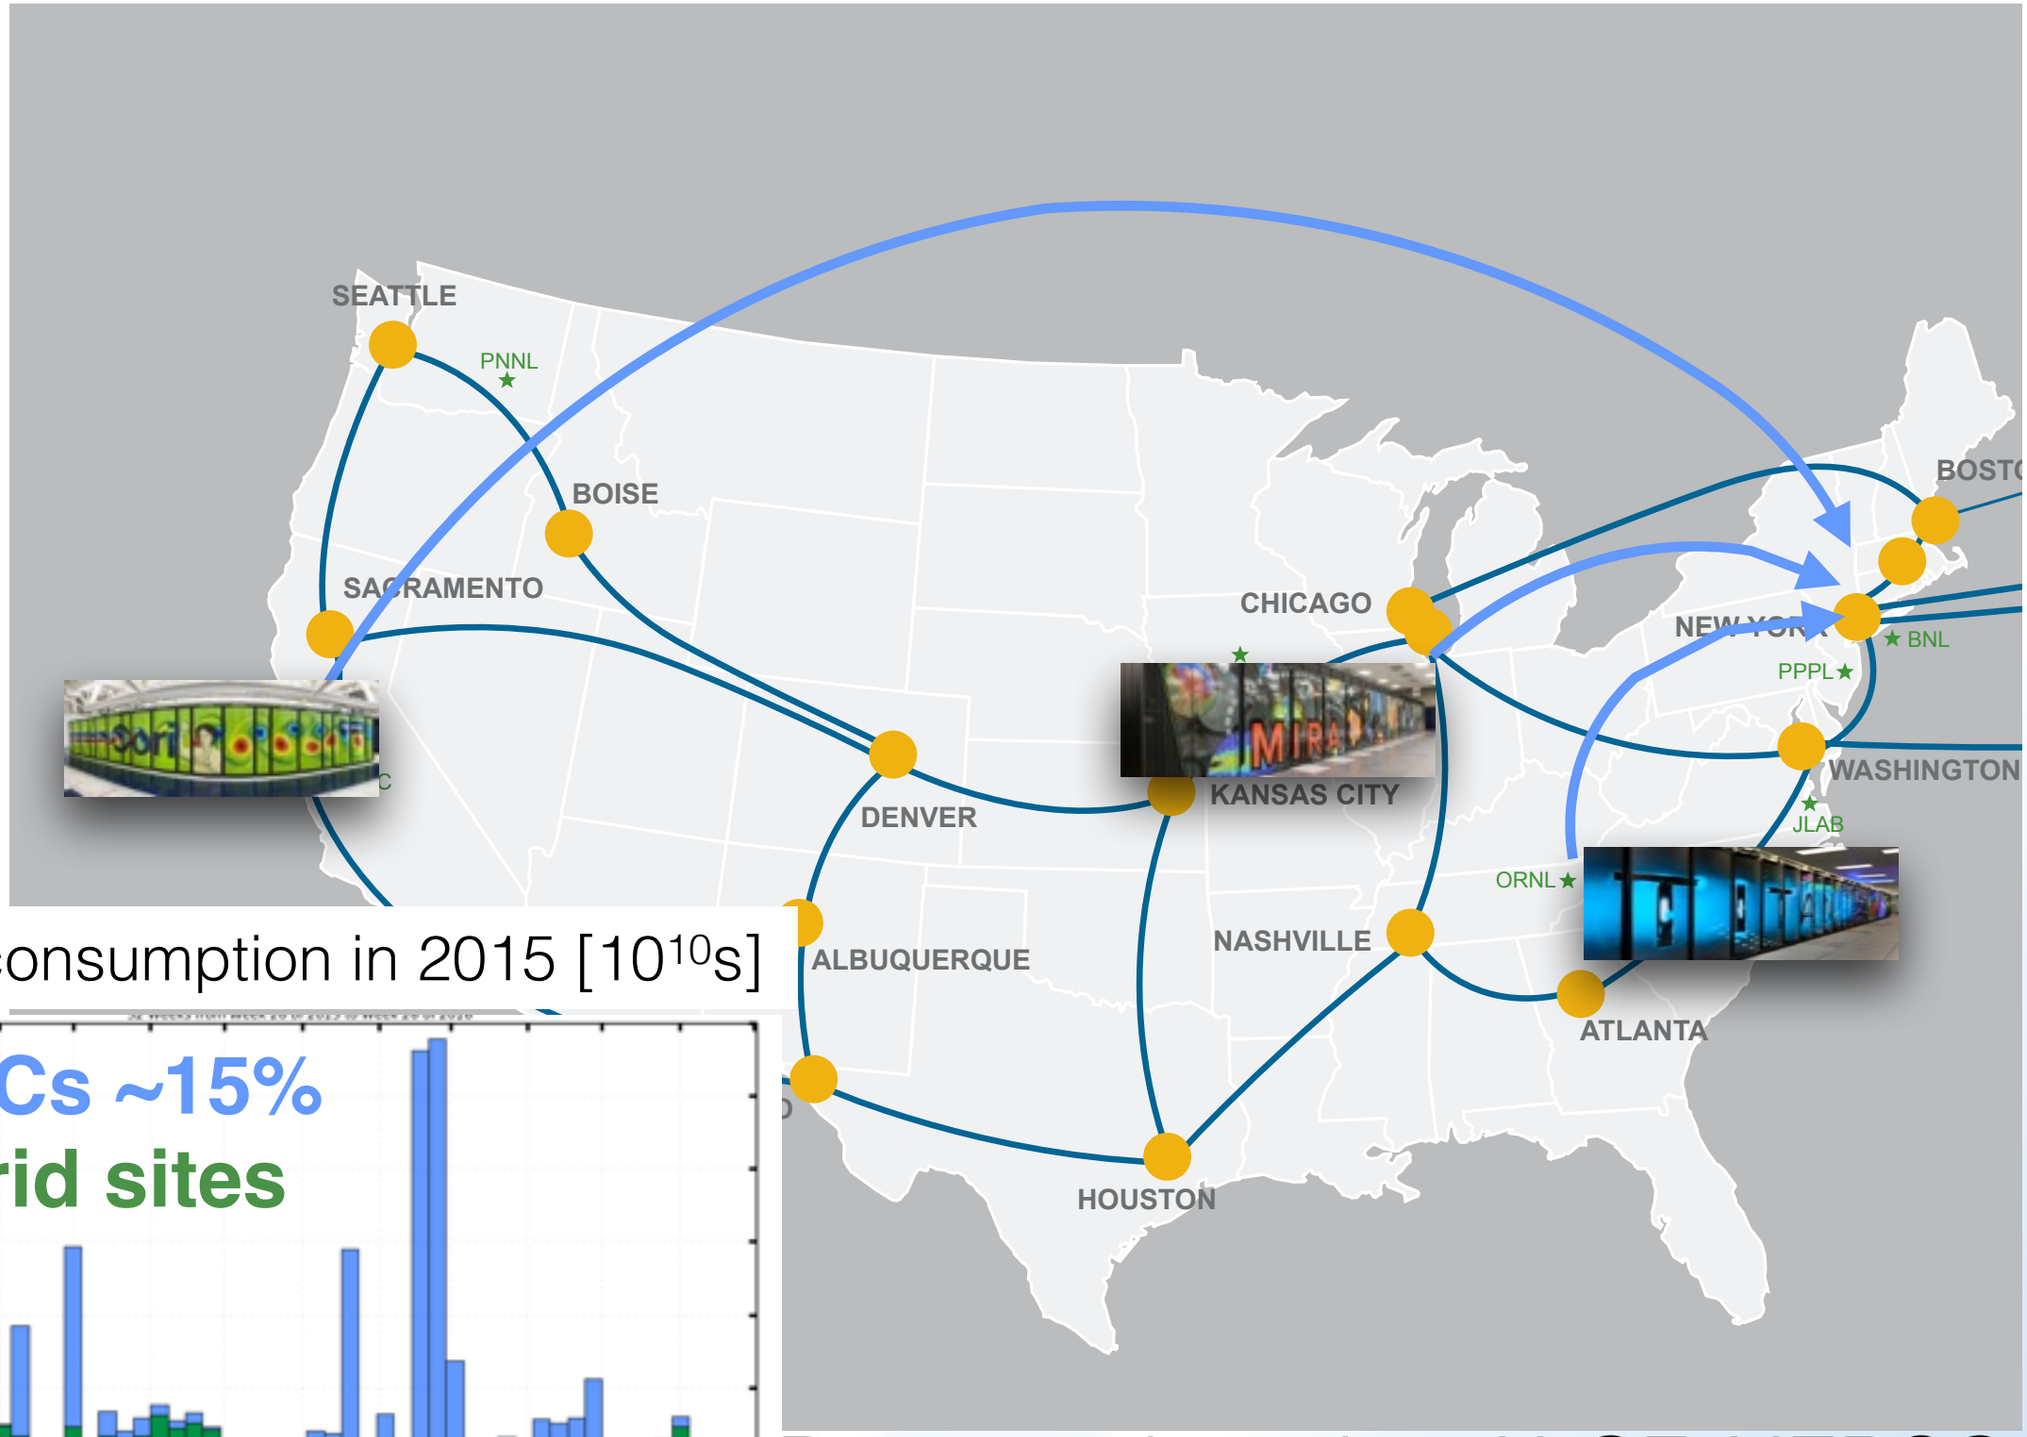

CPU consumption in 2015 [10^{10}s]

Data produced at ALCF, NERSC & OLCF stored at **BNL**

SDCC in 2017

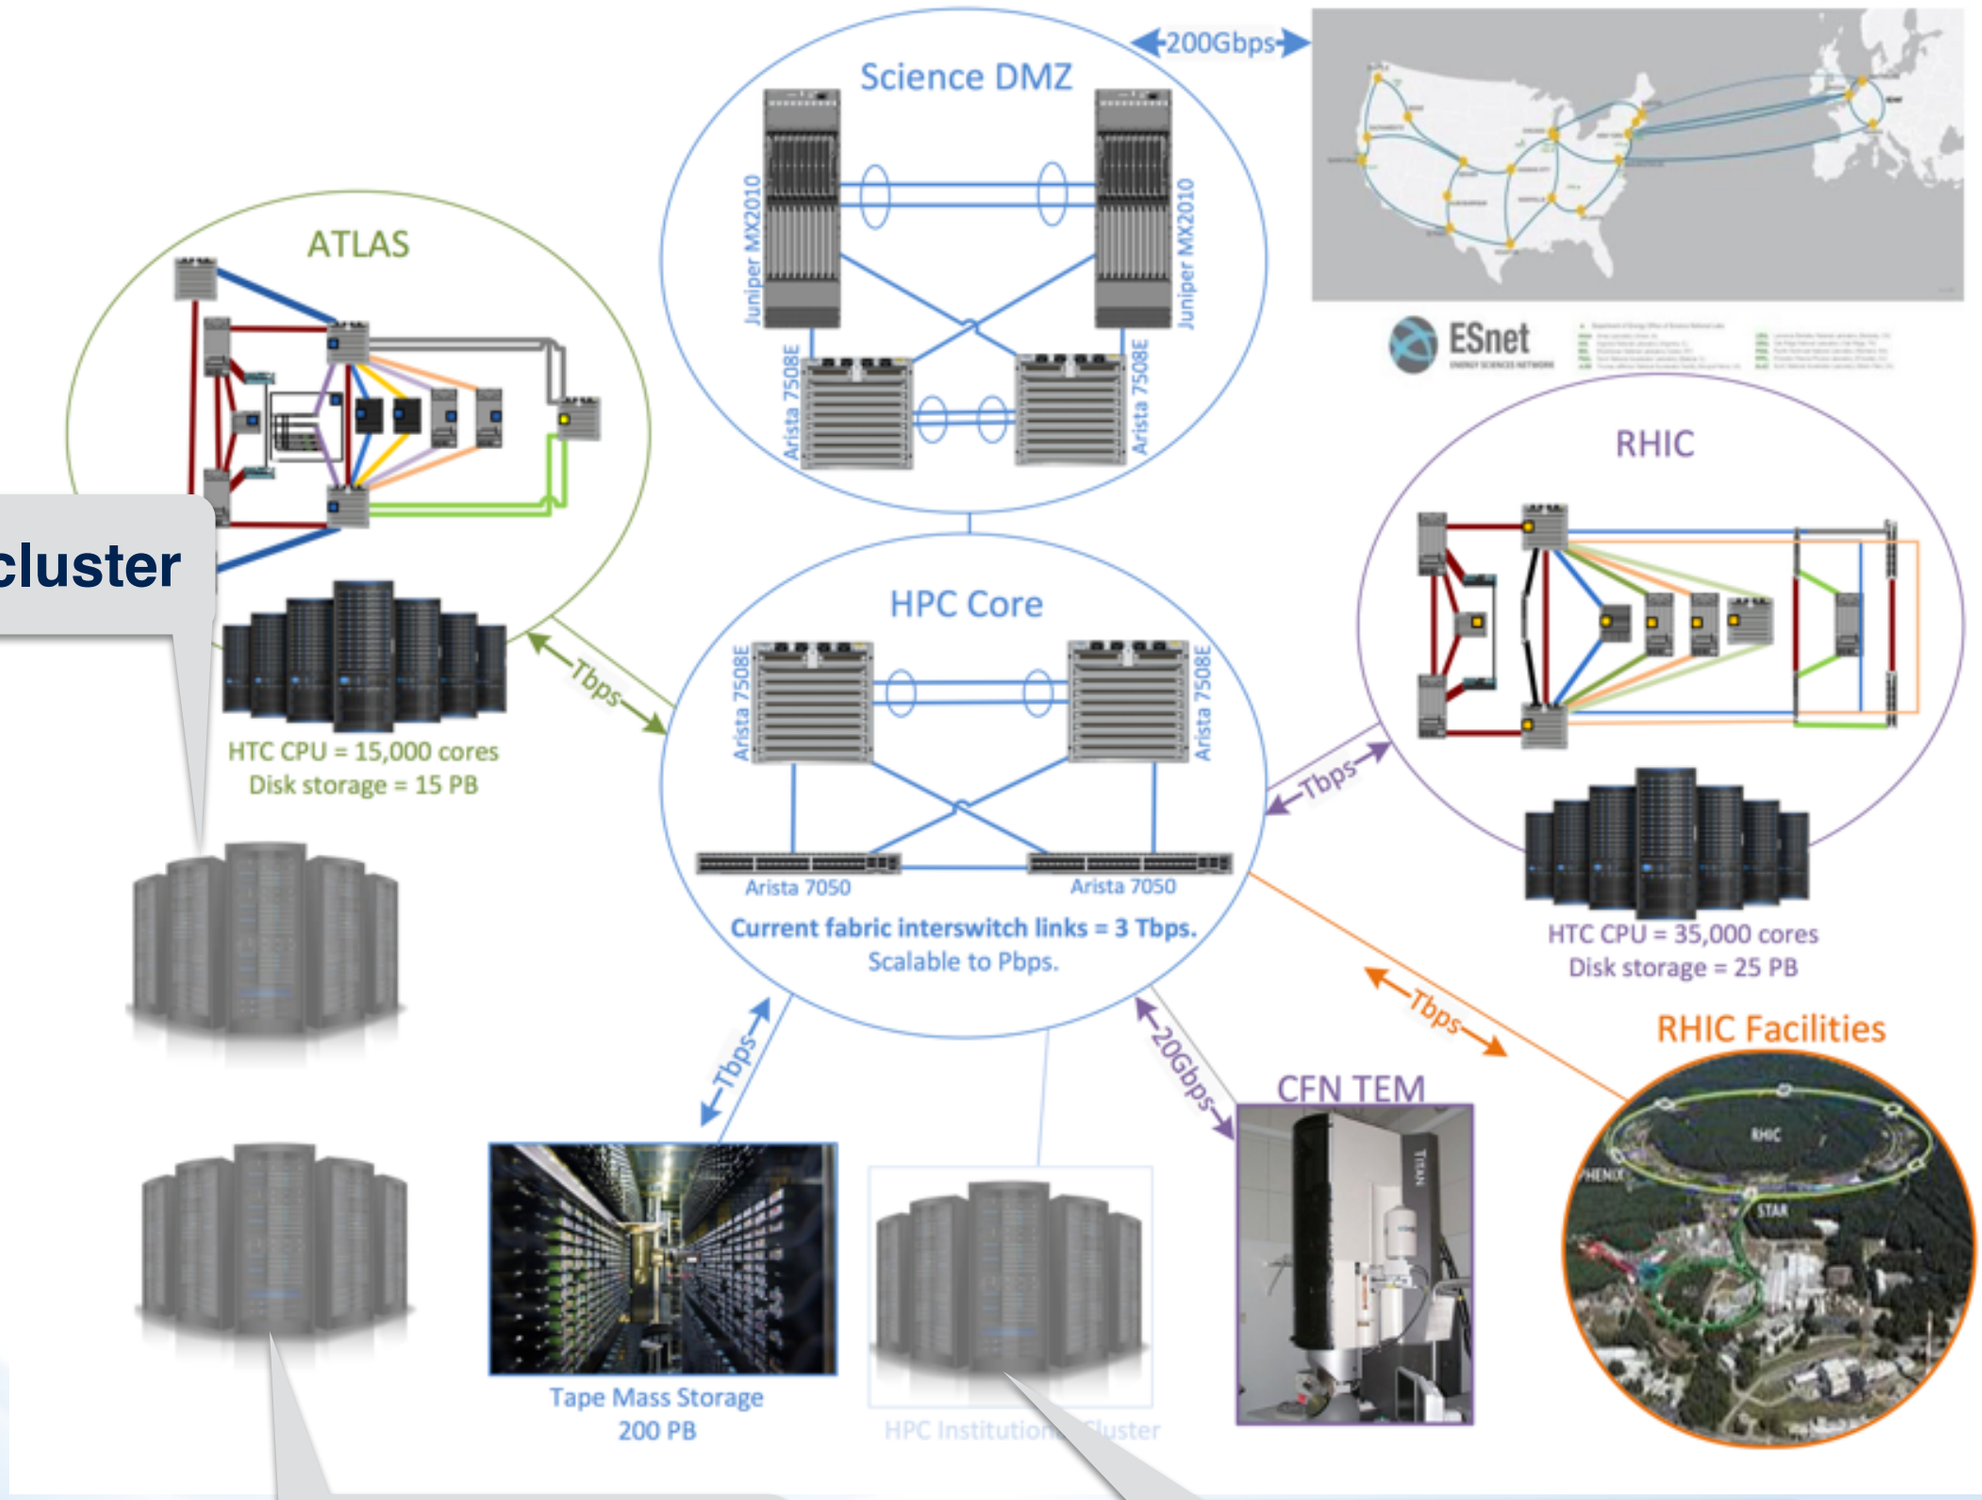

KNL cluster

USQCD cluster

Institutional Cluster

SDCC and RHIC - the next 5 years

- Resources for PHENIX & STAR experiments data reconstruction and analysis
  - Analysis continues more than 5 years after data taking
  - 1000s of users
- Simulation needs for
  - sPHENIX : detector optimisation and data taking
  - eRHIC optimisation

The Ongoing Science Mission

Status: RHIC-II configuration is complete

- Vertex detectors in STAR (HFT) and PHENIX
- Luminosity reaches 25x design luminosity

Plan: Complete RHIC mission in 3 campaigns:

- **2014–18: Heavy flavor probes of the QGP**
  Transverse spin physics
  Isobar system test of QCD anomalies
- *2018: Install low energy e-cooling*
- **2019/20: High precision scan of the QCD phase diagram & search for critical point**
- *2021: Install sPHENIX*
- **2022-??: Probe perfect liquid QGP with precision measurements of jet quenching and Upsilon suppression**
- *Transition to eRHIC?*

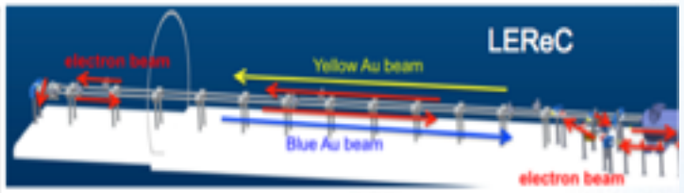

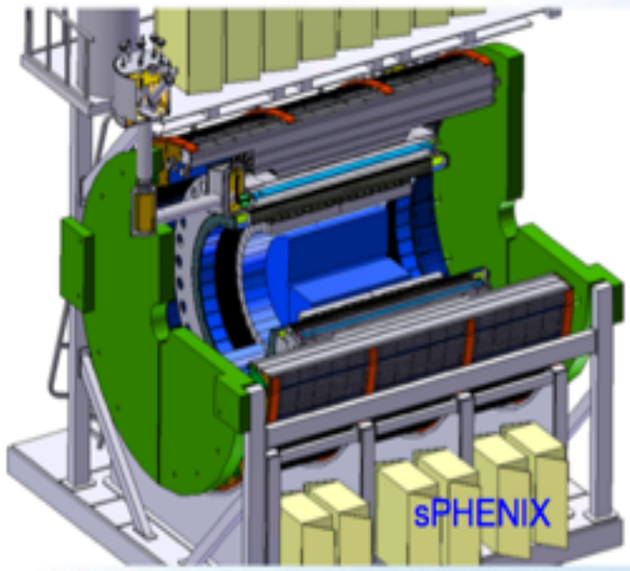

RHIC remains a unique discovery facility

Brookhaven Science Associates 2

BROOKHAVEN NATIONAL LABORATORY

ATLAS and the LHC Program

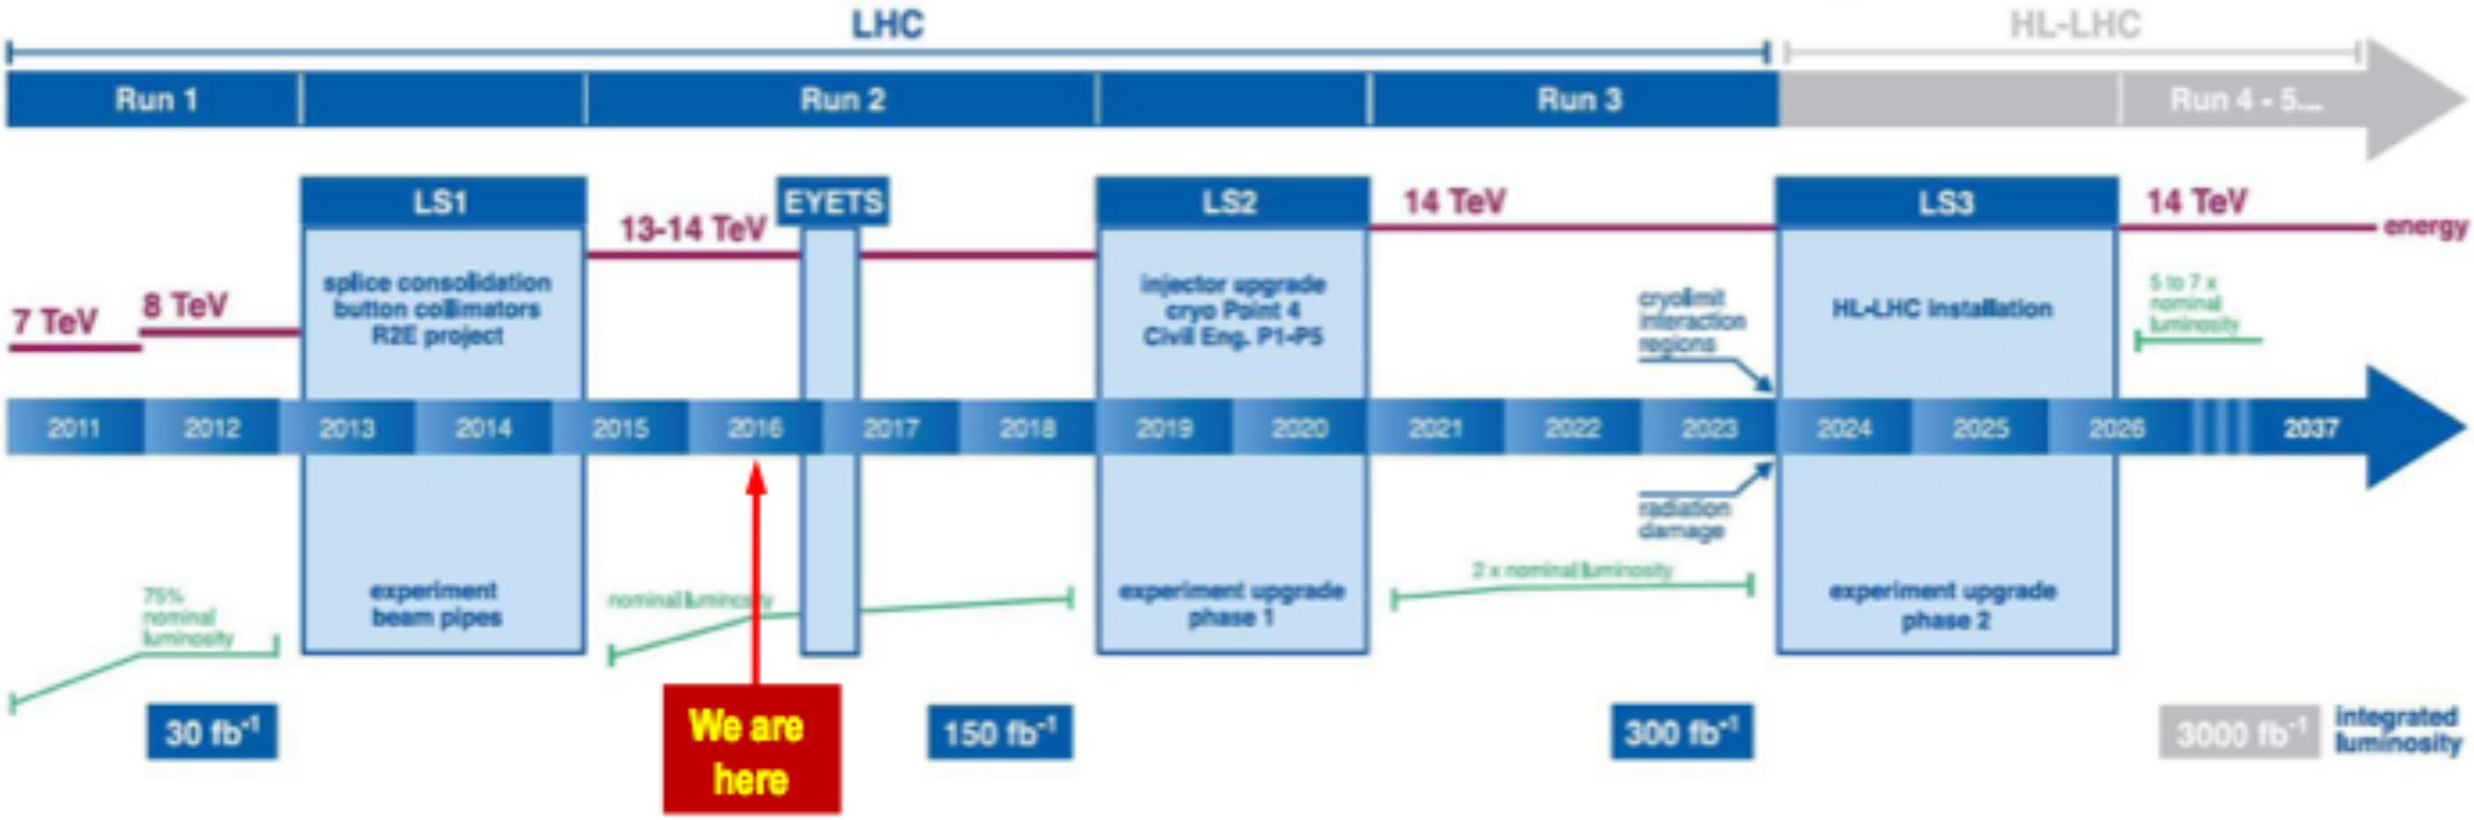

- High-Luminosity LHC (HL-LHC) project **Approved June 2016**
- A science program for the next 20 years
- **HL-LHC first priority of High Energy Physics program within the US**

Pile-up - the challenge for HL-LHC

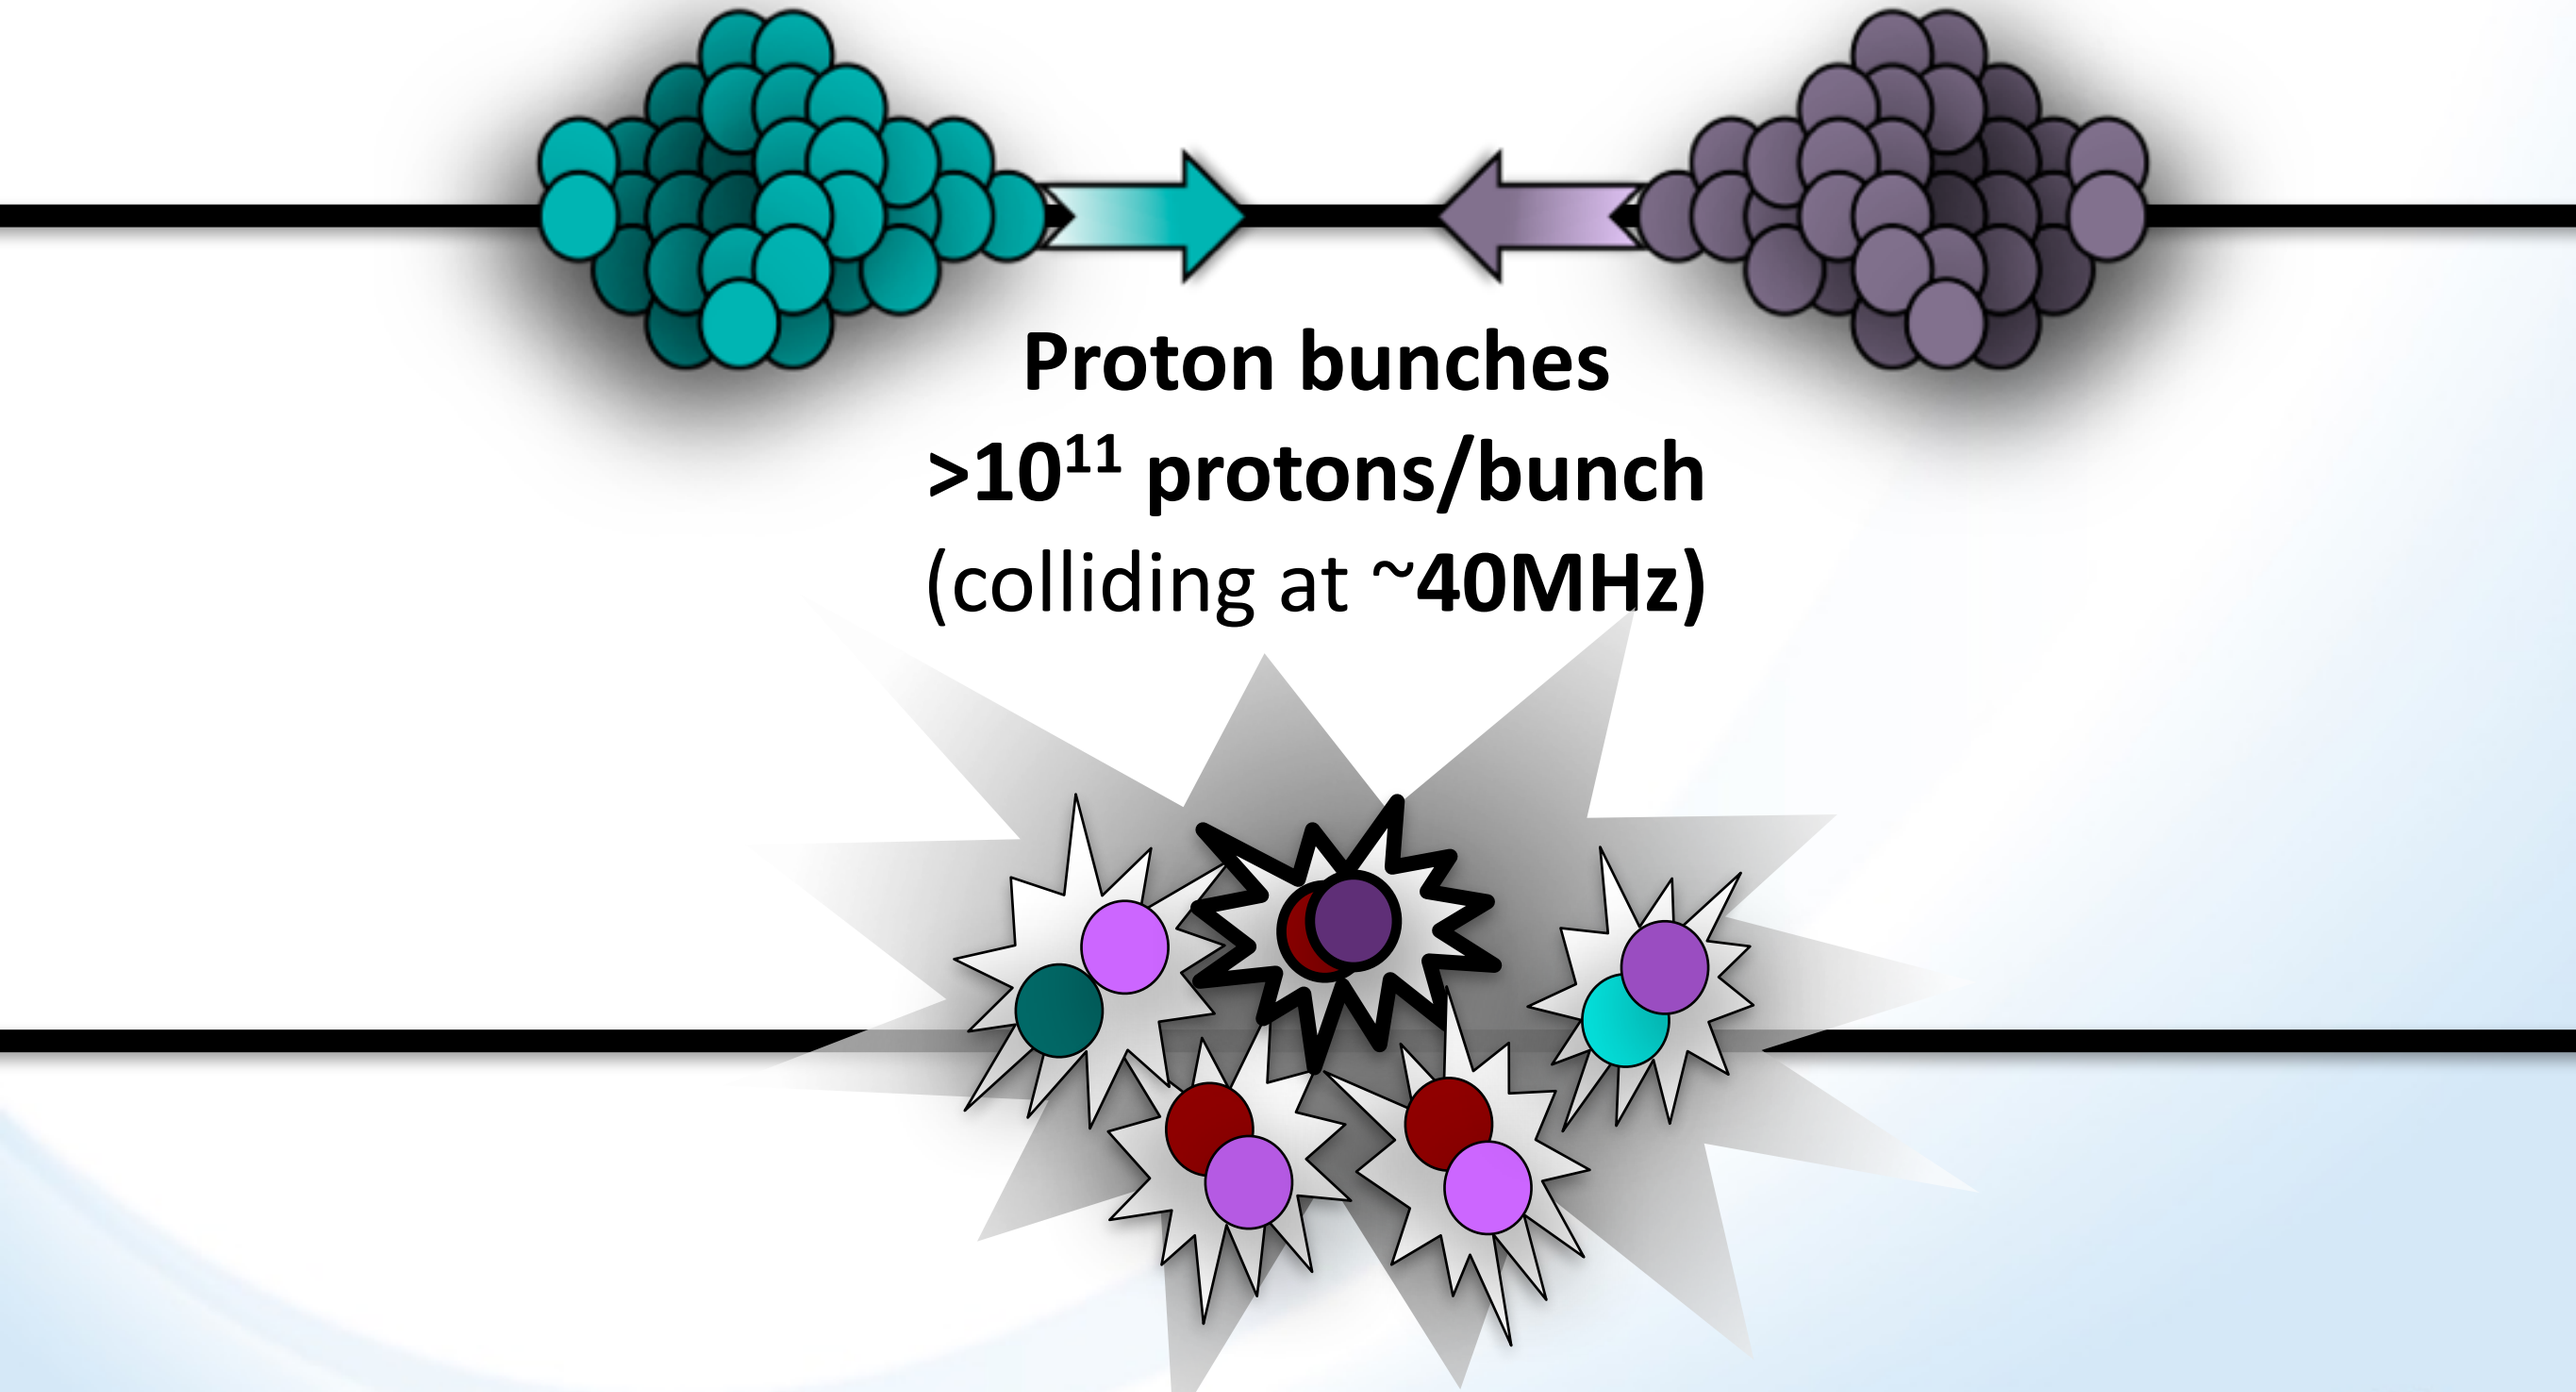

Proton bunches
>10^{11} protons/bunch
(colliding at ~**40MHz)**

N p-p collisions / bunch crossing

ATLAS and the LHC Program

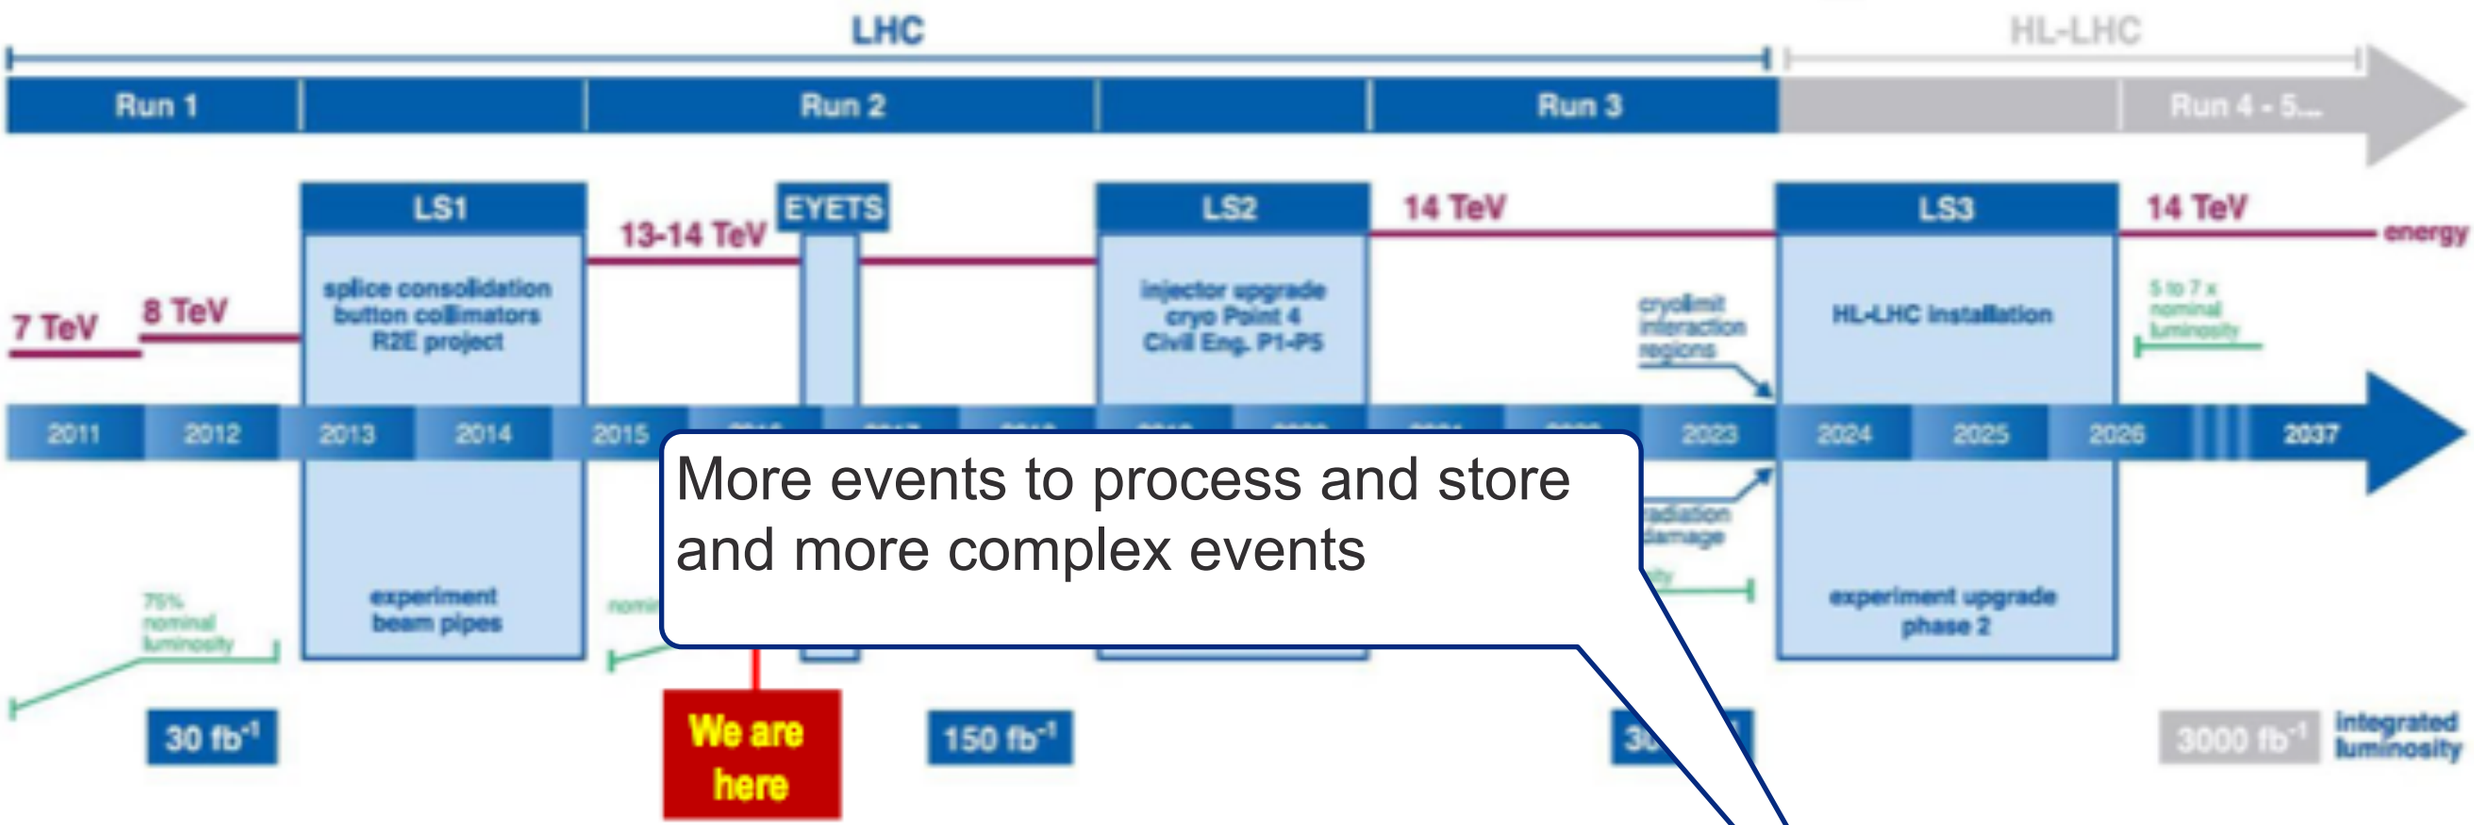

Trigger rate [kHz]	1	2	5-10
Number of events/ bunch crossing	20	40	150-200

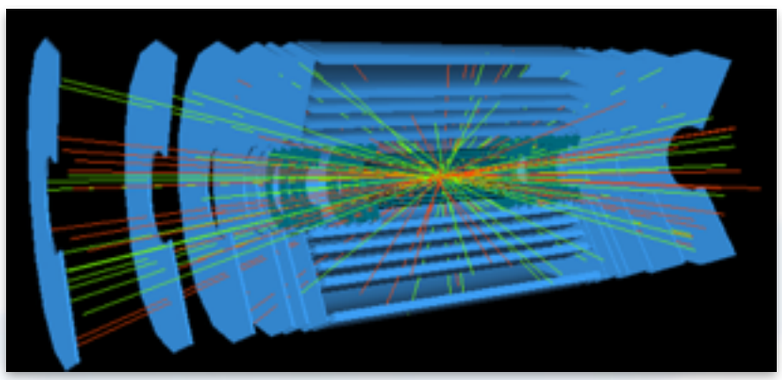

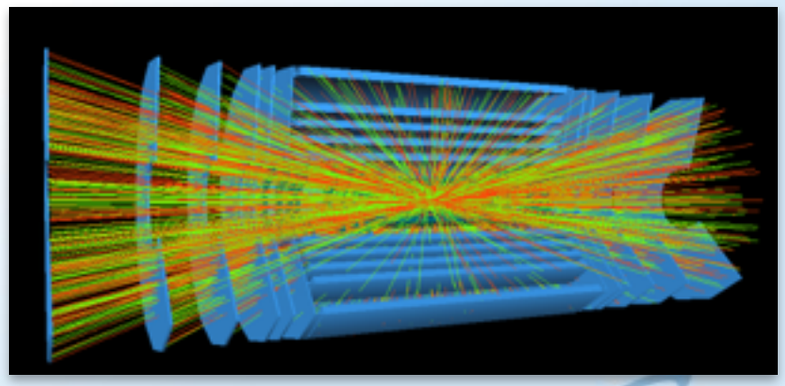

Computing Needs for HL-LHC

- Estimated computing needs beyond flat budget extrapolations :
  - Factor ~10 missing
  - Both for storage and computing

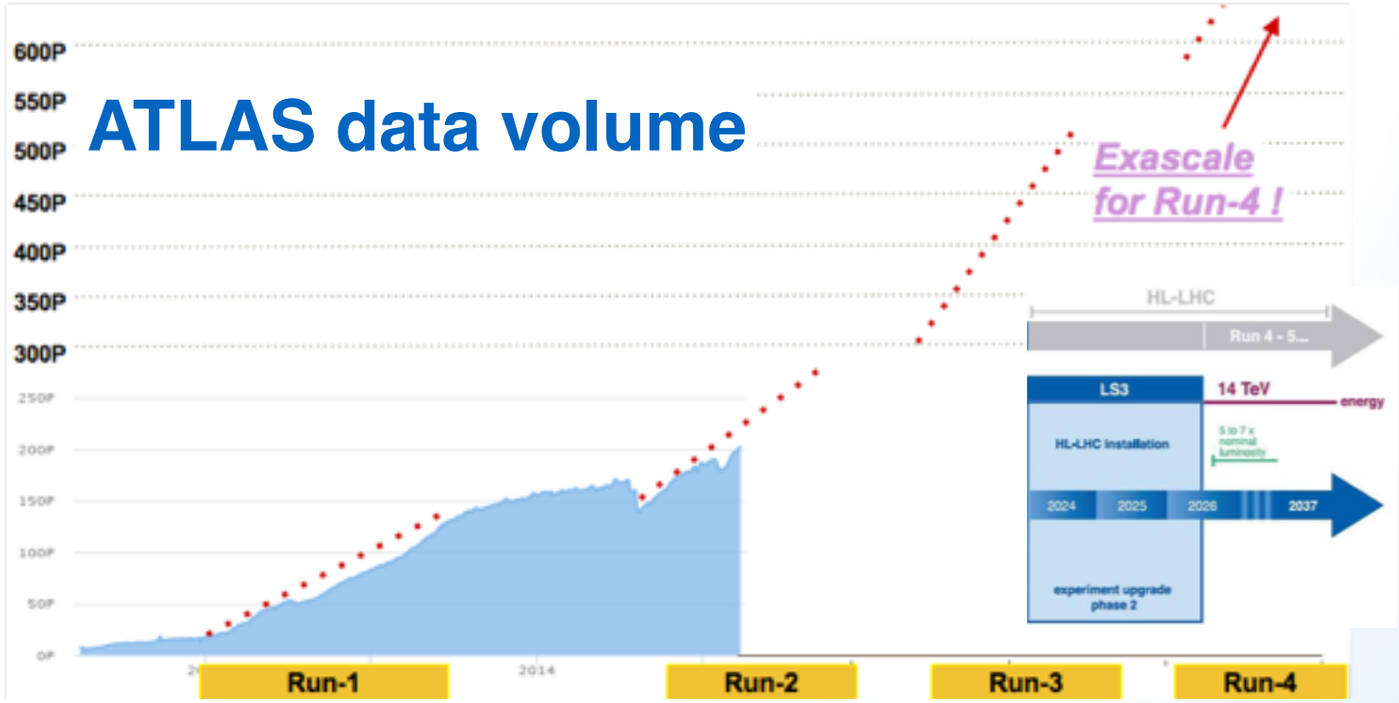

- A big challenge ahead
- US ATLAS Tier 1 at BNL has to host ~25% ATLAS computing capabilities

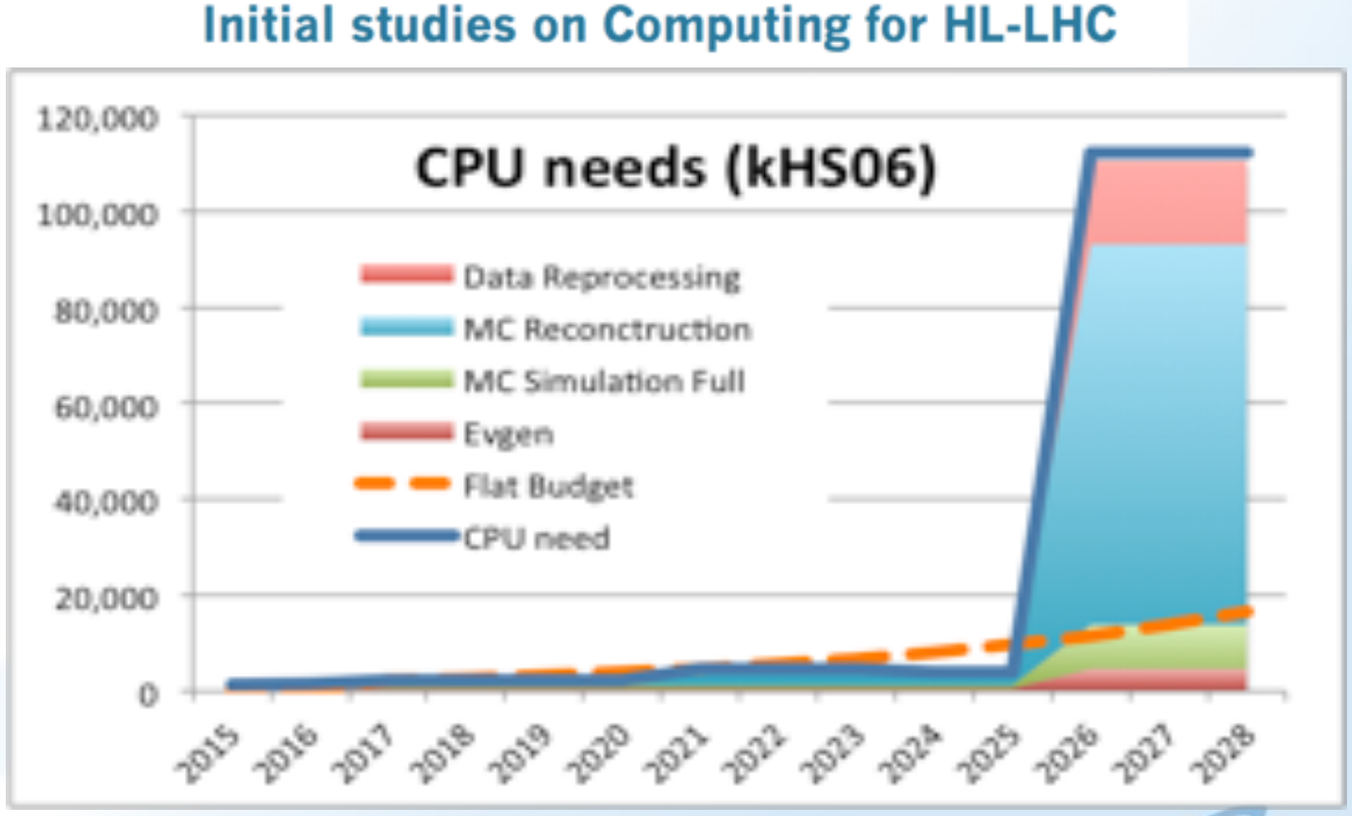

NSLS-II 10,000 times brighter than NSLS

- Data rates at NSLS-II will eventually be **> 90 TB/day** driven by use of high speed 2D detectors and the brightness of NSLS-II allowing for faster data collection.
- This results in total storage requirements of **~20 PB/y**, by 2020 being accessed by **3000+ users**.

- 21 more beamlines being developed in addition to the 7 construction project beamlines.

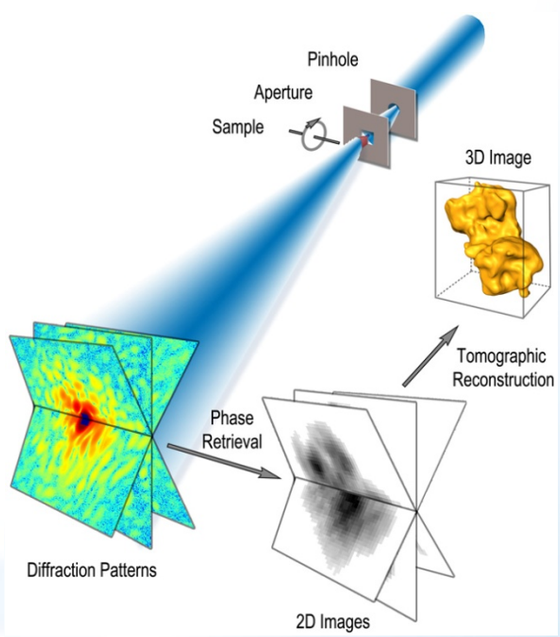

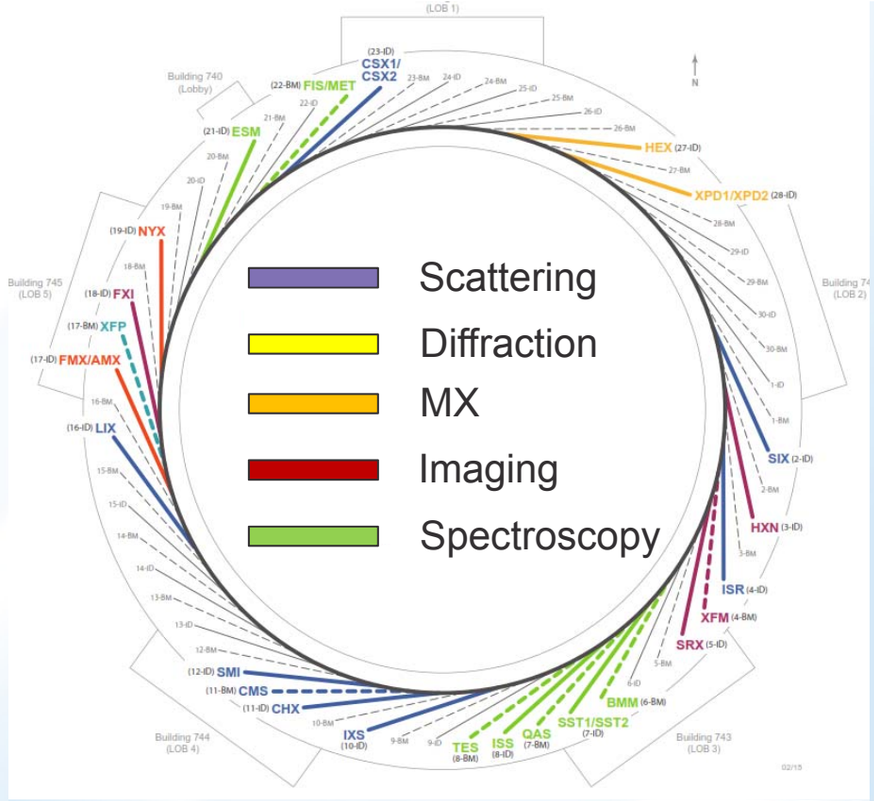

New computing centre – Conceptual Design

Delivery early 2021

60% Space Plan

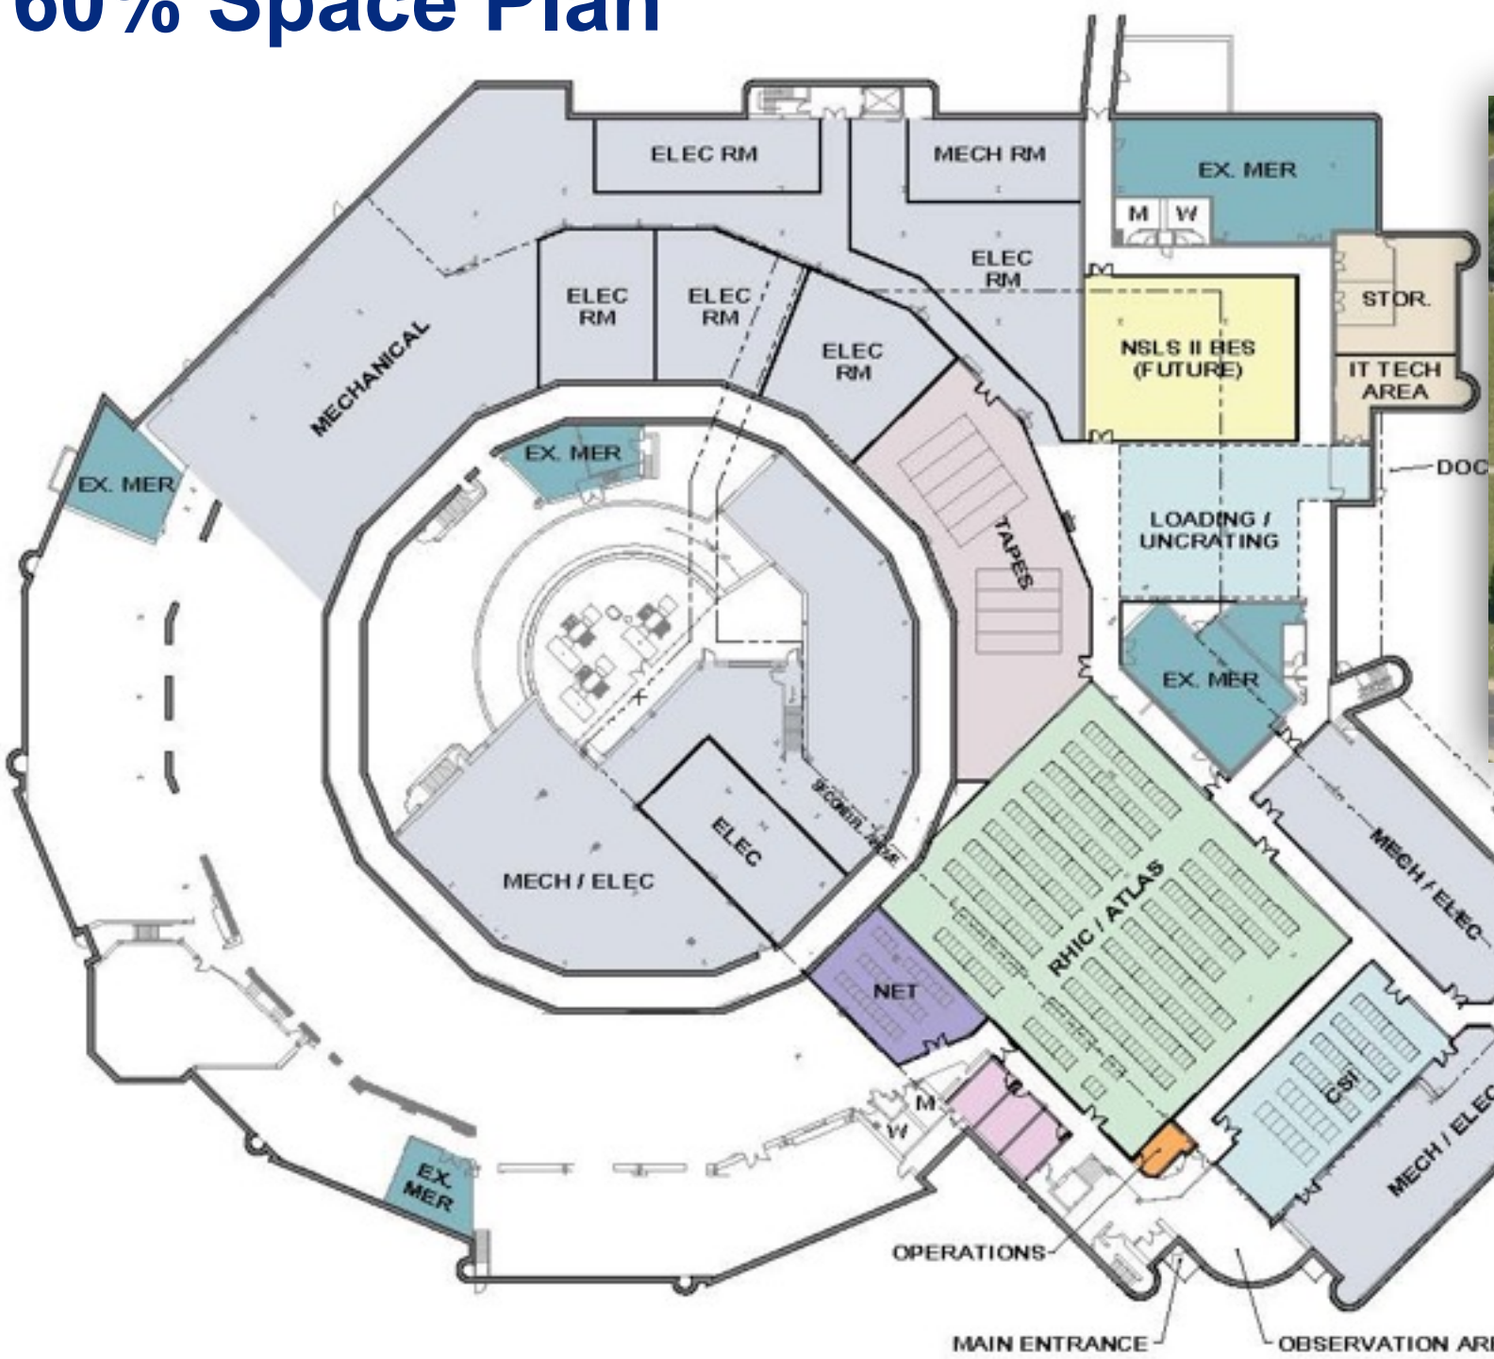

NSLS-I Building

~30k SqF

CD-1 review August 23-25

Summary and Outlook

- RACF successfully meet computational needs of US ATLAS and RHIC programs
  - a world class big-data processing facility
  - close to the users and responsive to their needs

- SDCC with CSI will leverage on this expertise
  - integration of National Nuclear Data Center, Atmospheric Radiation Measurement, NSLS-II,…

- Rapid growth of capacities adding HPC capabilities

- Challenging and exciting years in front of us!

Design – An Incremental Approach

- **Power**
  - Day-one capability (2021) – 2.4 MW IT power (dedicated computing power). This is approximately double current RACF IT power.
  - Provide provision for future1.2 MW IT power increments to 6MW Max.
- **Cooling**
  - Day-one cooling capability to support 2.4 MW IT power
  - Provide provision for future 1.2 MW IT power deployments
- **Space**
  - Provide opportunity for future (long term) growth. Both computing and offices.